Enumerating k-Vertex Connected Components in Large Graphs

Dong Wen[♮], Lu Qin[♮], Xuemin Lin[♯], Ying Zhang[♮], and Lijun Chang[♯]
[♮]CAI, University of Technology, Sydney, Australia
[♯]The University of New South Wales, Australia
[♮]dong.wen@student.uts.edu.au; {lu.qin, ying.zhang}@uts.edu.au;
[♯]{lxue, ljchang}@cse.unsw.edu.au;

arXiv:1703.08668v1 [cs.DB] 25 Mar 2017

ABSTRACT

Cohesive subgraph detection is an important graph problem that is widely applied in many application domains, such as social community detection, network visualization, and network topology analysis. Most of existing cohesive subgraph metrics can guarantee good structural properties but may cause the free-rider effect. Here, by free-rider effect, we mean that some irrelevant subgraphs are combined as one subgraph if they only share a small number of vertices and edges. In this paper, we study k-vertex connected component (k-VCC) which can effectively eliminate the free-rider effect but less studied in the literature. A k-VCC is a connected subgraph in which the removal of any $k-1$ vertices will not disconnect the subgraph. In addition to eliminating the free-rider effect, k-VCC also has other advantages such as bounded diameter, high cohesiveness, bounded graph overlapping, and bounded subgraph number. We propose a polynomial time algorithm to enumerate all k-VCCs of a graph by recursively partitioning the graph into overlapped subgraphs. We find that the key to improving the algorithm is reducing the number of local connectivity testings. Therefore, we propose two effective optimization strategies, namely neighbor sweep and group sweep, to largely reduce the number of local connectivity testings. We conduct extensive performance studies using seven large real datasets to demonstrate the effectiveness of this model as well as the efficiency of our proposed algorithms.

1. INTRODUCTION

Graphs have been widely used to represent the relationships of entities in the real world. With the proliferation of graph applications, research efforts have been devoted to many fundamental problems in mining and analyzing graph data. Recently, cohesive subgraph detection has drawn intense research interest [22]. Such problem can be widely adopted in many real-world applications, such as community detection [10, 16], network clustering [26], graph visualization [1, 35], protein-protein network analysis [2], and system analysis [34].

In the literature, a large number of cohesive subgraph models have been proposed. Among them, a clique, in which every pair of vertices are connected, guarantees perfect familiarity and reachability among vertices. Since the definition of the clique is too strict, clique-relaxation models are proposed in the literature including s-clique [19], s-club [19], γ-quasi-clique [33] and k-plex [4, 23]. Nevertheless, these models require exponential computation time and may lack guaranteed cohesiveness. To conquer this problem, other models are proposed such as k-core [3], k-truss [9, 27, 24], k-mutual-friend subgraph [36] and k-ECC (k-edge connected component) [37, 6], which require polynomial computation time and guarantee decent cohesiveness. For example, a k-core guarantees that every vertex has a degree at least k in the subgraph, and a k-ECC guarantees that the subgraph cannot be disconnected after removing any $k-1$ edges.

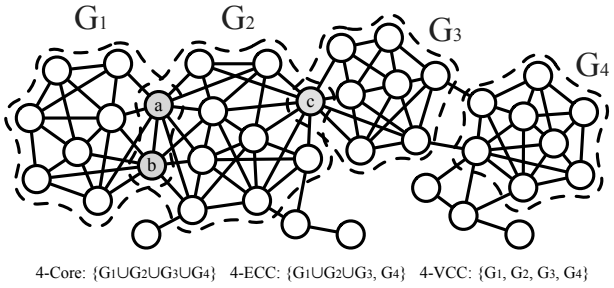

Figure 1: Cohesive subgraphs in graph G.

Motivation. Despite the good structural guarantees in existing cohesive subgraph models, we find that most of these models cannot effectively eliminate the free-rider effect. Here, by free-rider effect, we mean that some irrelevant subgraphs are combined as one result subgraph if they only share a small number of vertices and edges. To illustrate the free rider effect, we consider a graph G shown in Fig. 1, which includes four subgraphs G_1, G_2, G_3, and G_4. The four subgraphs are loosely connected because: G_1 and G_2 share a single edge (a, b); G_2 and G_3 share a single vertex c; and G_3 and G_4 do not share any edge or vertex. Let $k = 4$. Based on the k-core model, there is only one k-core, which is the union of the four subgraphs G_1, G_2, G_3, and G_4, along with the two edges connecting G_3 and G_4. Based on the k-ECC model, there are two k-ECCs, which are G_4 and the union of three subgraphs G_1, G_3, and G_3. Motivated by this, we aim to detect cohesive subgraphs and effectively eliminate the free-rider effect, i.e., to accurately detect G_1, G_2, G_3 and G_4 as result cohesive subgraphs in Fig. 1.

In the literature, a recent work [31] aims to eliminate the free-rider effect in local community search. Given a query vertex, the algorithm in [31] tries to eliminate the free-rider effect by weighting each vertex in the graph by its proximity to the query vertex. Based on the vertex weights, a query-biased subgraph is returned by considering both the density and the proximity to the query ver-

tex. Unfortunately, such a query-biased local community model cannot be used in cohesive subgraph detection.

k-Vertex Connected Component. Vertex connectivity, which is also named structural cohesion [20], is *the minimum number of vertices that need to be removed to disconnect the graph*. It has been proved as an outstanding metric to evaluate the cohesiveness of a social group [20, 29]. We find this sociological conception can be used to detect cohesive subgraphs and effectively eliminate the free-rider effect. Given an integer k, a k-vertex connected component (k-VCC) is *a maximal connected subgraph in which the removal of any $k-1$ vertices cannot disconnect the subgraph*. Given a graph G and a parameter k, we aim to detect all k-VCCs in G. In Fig. 1 and $k = 4$, there are four k-VCCs G_1, G_2, G_3, and G_4 in G. The subgraph formed by the union of G_1 and G_2 is not a k-VCC because it will be disconnected by removing two vertices a and b.

Effectiveness. k-VCC effectively eliminates the free-rider effect by ensuring that each k-VCC cannot be disconnected by removing any $k-1$ vertices. In addition, k-VCC also have the following four good structural properties.

- *Bounded Diameter.* The diameter of a k-VCC $G'(V', E')$ is bounded by $\lfloor \frac{|V'|-2}{\kappa(G')} \rfloor + 1$ where $\kappa(G')$ is the vertex connectivity of G'. For example, we consider the 4-VCC G_1 with 9 vertices in Fig. 1. The diameter of G_1 is bounded by 2.
- *High Cohesiveness.* We can guarantee that a k-VCC is nested in a k-ECC and a k-core. Therefore, a k-VCC is generally more cohesive and inherits all the structural properties of a k-core and a k-ECC. For example, each of the four 4-VCCs in Fig. 1 is also a 4-core and a 4-ECC.
- *Subgraph Overlapping.* Unlike k-core and k-ECC, k-VCC model allows overlapping between k-VCCs, and we can guarantee that the number of overlapped vertices for any pair of k-VCCs is smaller than k. For example, the two 4-VCCs G_1 and G_2 in Fig. 1 overlap two vertices and an edge.
- *Bounded Subgraph Number.* Even with overlapping, we can bound the number of k-VCCs to be linear to the number of vertices in the graph. This indicates that redundancies in the k-VCCs are limited. For example, the graph shown in Fig. 1 contains four 4-VCCs with three vertices a, b, and c duplicated.

The details of the four properties can be found in Section 2.2.

Efficiency. In this paper, we propose an algorithm to enumerate all k-VCCs in a given graph G via overlapped graph partition. Briefly speaking, we aim to find a vertex cut with fewer than k vertices in G. Here, a vertex cut of G is a set of vertices the removal of which disconnects the graph. With the vertex cut, we can partition G into overlapped subgraphs each of which contains all the vertices in the cut along with their induced edges. We recursively partition each of the subgraphs until no such cut exists. In this way, we compute all k-VCCs. For example, suppose the graph G is the union of G_1 and G_2 in Fig. 1, $k = 4$, we can find a vertex cut with two vertices a and b. Thus we partition the graph into two subgraphs G_1 and G_2 that overlap two vertices a, b and an edge (a, b). Since neither G_1 nor G_2 has any vertex cut with fewer then k vertices, we return G_1 and G_2 as the final k-VCCs. We theoretically analyze our algorithm and prove that *the set of k-VCCs can be enumerated in polynomial time*. More details can be found in Section 4.3.

Nevertheless, the above algorithm has a large improvement space. The most crucial operation in the algorithm is called local connectivity testing, which given two vertices u and v, tests whether u and v can be disconnected in two components by removing at most $k-1$ vertices from G. To find a vertex cut with fewer than k vertices, we need to conduct local connectivity testing between a source vertex s and each of other vertices v in G in the worst case. Therefore, *the key to improving algorithmic efficiency is to reduce the number of local connectivity testings in a graph.* Given a source vertex s, if we can avoid testing the local connectivity between s and a certain vertex v, we call it as we can *sweep vertex v*. We propose two strategies to sweep vertices.

- *Neighbor Sweep.* If a vertex has certain properties, all its neighbors can be swept. Therefore, we call this strategy neighbor sweep. Moreover, we maintain a deposit value for each vertex, and once we finish testing or sweep a vertex, we increase the deposit values for its neighbors. If the deposit value of a vertex satisfies certain condition, such vertex can also be swept.
- *Group Sweep.* We introduce a method to divide vertices in a graph into disjoint groups. If a vertex in a group has certain properties, vertices in the whole group can be swept. We call this strategy group sweep. Moreover, we maintain a group deposit value for each group. Once we test or sweep a vertex in the group, we increase the corresponding group deposit value. If the group deposit value satisfies certain conditions, vertices in such whole group can also be swept.

Even though these two strategies are studied independently, they can be used together and boost the effectiveness of each other. With these two vertex sweep strategies, we can significantly reduce the number of local connectivity testings in the algorithm. Experimental results show the excellent performance of our sweep strategies. More details can be found in Section 5 and Section 6.

Contributions. We make the following contributions in this paper.

(1) Theoretical analysis for the effectiveness of k-VCC. We present several properties to show the excellent quality of k-vertex connected component. Although the concept of vertex connectivity has been studied in the literature to evaluate the cohesiveness of a social group, this is the first work that aims to enumerate all k-VCCs and considers free-rider effect elimination in cohesive subgraph detection to the best of our knowledge.

(2) A polynomial time algorithm based on overlapped graph partition. We propose an algorithm to compute all k-VCCs in a graph G. The algorithm recursively divides the graph into overlapped subgraphs until each subgraph cannot be further divided. We prove that our algorithm terminates in polynomial time.

(3) Two effective pruning strategies. We design two pruning strategies, namely neighbor sweep and group sweep, to largely reduce the number of local connectivity testings and thus significantly speed up the algorithm.

(4) Extensive performance studies. We conduct extensive performance studies on 7 real large graphs to demonstrate the effectiveness of k-VCC and the efficiency of our proposed algorithms.

Outline. The rest of this paper is organized as follows. Section 2 formally defines the problem and presents its rationale. Section 3 gives a framework to compute all k-VCCs in a given graph. Section 4 gives a basic implementation of the framework and analyzes the time complexity of the algorithm. Section 5 introduces several strategies to speed up the algorithm. Section 6 evaluates the model and algorithms using extensive experiments. Section 7 reviews related works and Section 8 concludes the paper.

2. PRELIMINARY

2.1 Problem Statement

In this paper, we consider an undirected and unweighted graph $G(V, E)$, where V is the set of vertices and E is the set of edges. We also use $V(G)$ and $E(G)$ to denote the set of vertices and edges of graph G respectively. The number of vertices and the number of edges are denoted by $n = |V|$ and $m = |E|$ respectively. For simplicity and without loss of generality, we assume that G is a connected graph. We denote neighbor set of a vertex u by $N(u)$, i.e., $N(u) = \{u \in V | (u, v) \in E\}$, and degree of u by $d(u) = |N(u)|$. Given two graphs g and g', we use $g \subseteq g'$ to denote that g is a subgraph of g'. Given a set of vertices V_s, the induced subgraph $G[V_s]$ is a subgraph of G such that $G[V_s] = (V_s, \{(u, v) \in E | u, v \in V_s\})$. For any two subgraphs g and g' of G, we use $g \cup g'$ to denote the union of g and g', i.e., $g \cup g' = (V(g) \cup V(g'), E(g) \cup E(g'))$. Before stating the problem, we firstly give some basic definitions.

DEFINITION 1. (VERTEX CONNECTIVITY) *The vertex connectivity of a graph G, denoted by $\kappa(G)$, is defined as the minimum number of vertices whose removal results in either a disconnected graph or a trivial graph (a single-vertex graph).*

DEFINITION 2. (K-VERTEX CONNECTED) *A graph G is k-vertex connected if: 1) $|V(G)| > k$; and 2) the remaining graph is still connected after removing any $(k-1)$ vertices. That is, $\kappa(G) \geq k$.*

We use the term *k-connected* for short when the context is clear. It is easy to see that any nontrivial connected graph is at least 1-connected. Based on Definiton 2, we define the *k-Vertex Connected Component* (k-VCC) as follows.

DEFINITION 3. (K-VERTEX CONNECTED COMPONENT) *Given a graph G, a subgraph g is a k-vertex connected component (k-VCC) of G if: 1) g is k-vertex connected; and 2) g is maximal. That is, $\nexists g' \subseteq G$, such that $\kappa(g') \geq k$, $g \subseteq g'$.*

Problem Definition. Given a graph G and an integer k, we denote the set of all k-VCCs of G as $VCC_k(G)$. In this paper, we study the problem of efficiently enumerating all k-VCCs of G, i.e, to compute $VCC_k(G)$.

EXAMPLE 1. *For the graph G in Fig. 1, given parameter $k = 4$, there are four 4-VCCs: $VCC_4(G) = \{G_1, G_2, G_3, G_4\}$. We cannot disconnect each of them by removing any 3 or fewer vertices. Subgraph $G_1 \cup G_2$ is not a 4-VCC because it will be disconnected after removing two vertices a and b.*

2.2 Why k-Vertex Connected Component?

k-VCC model effectively reduces the free-rider effect by ensuring that each k-VCC cannot be disconnected by removing any $k-1$ vertices. In this subsection, we show other good structural properties of k-VCC in terms of bounded diameter, high cohesiveness, bounded overlapping and bounded component number. None of other cohesive graph models, such as k-core and k-Edge Connected Component (k-ECC) can achieve these four goals simultaneously.

Diameter. Before discussing the diameter of a k-VCC, we first quote Global Menger's Theorem as follows.

THEOREM 1. *A graph is k-connected if and only if any pair of vertices u,v is joined by at least k vertex-independent u-v paths. [18]*

This theorem shows the equivalence of vertex connectivity and the number of vertex-independent paths, both of which are considered as important properties for graph cohesion [29]. Based on this theorem, we can bound the diameter of a k-VCC, where the diameter of a graph G, denoted by $diam(G)$, is the longest shortest path between any pair of vertices in G:

$$diam(G) = max_{u,v \in V(G)} dist(u, v, G) \quad (1)$$

Here, $dist(u, v, G)$ is the shortest distance of the pair of vertices u and v in G. Small diameter is considered as an important feature for a good community in [11]. We give the diameter upper bound for a k-VCC as follows.

THEOREM 2. *Given any k-VCC G_i of G, we have:*

$$diam(G_i) \leq \lfloor \frac{|V(G_i)| - 2}{\kappa(G_i)} \rfloor + 1. \quad (2)$$

PROOF. Consider any two vertices u and v in G_i, we have $d(u, v, G_i) \leq diam(G_i)$. Theorem 1 indicates that there exist at least $\kappa(G_i)$ vertex-disjoint paths between u and v in G_i, and in each path, we have at most $diam(G_i) - 1$ internal vertices since $dist(u, v, G_i) \leq diam(G_i)$. Thus we have at most $\kappa(G_i) \times (diam(G_i) - 1)$ internal vertices between them. With two endpoints u and v, we have $2 + \kappa(G_i) \times (diam(G_i) - 1) \leq |V(G_i)|$. Thus the upper bound of $diam(G_i)$ is $\lfloor \frac{|V(G_i)|-2}{\kappa(G_i)} \rfloor + 1$. □

Cohesiveness. To further investigate the quality of k-VCC, we introduce the Whitney Theorem [30]. Given a graph g, it analyzes the inclusion relation between vertex connectivity $\kappa(g)$, edge connectivity $\kappa'(g)$ and minimum degree $\delta(g)$. The theorem is presented as follows.

THEOREM 3. *For any graph g, $\kappa(g) \leq \kappa'(g) \leq \delta(g)$.*

From this theorem, we know that for a graph G, every k-VCC of G is nested in a k-ECC in G, and every k-ECC of G is nested in a k-core in G. Therefore, k-VCC is generally more cohesive than k-ECC and k-core.

Overlapping. The k-VCC model also supports vertex overlap between different k-VCCs, which is especially important in social networks. We can easily deduce the following property from the definition of k-VCC to bound the overlapping size.

PROPERTY 1. *Given two k-VCCs G_i and G_j in graph G, the number of overlapped vertices of G_i and G_j is less than k. That is, $|V(G_i) \cap V(G_j)| < k$.*

EXAMPLE 2. *In Fig. 1, we find two vertices a and b that are contained in two 4-VCCs G_1 and G_2. For k-ECC and k-core, two components will be combined if they have one vertex in common. For example, there is only one 4-core, which is the union of G_1, G_2, G_3 and G_4.*

Component Number. Once we allow overlapping between different components, the number of components can hardly be bounded. For example, the number of maximal cliques achieves $3^{\frac{n}{3}}$ for a graph with n vertices [21]. Nevertheless, we find that the number of k-VCCs in a graph G can be bounded by a function that is linear to the number of vertices in G. That is:

$$|VCC_k(G)| \leq \frac{|V(G)|}{2} \quad (3)$$

Detailed discussions of Eq. 3 can be found in Section 4. It is worth noting that the linear number of k-VCCs allows us to design a polynomial time algorithm to enumerate all k-VCCs. We will also discuss this in detail in Section 4.

Algorithm 1 KVCC-ENUM(G, k)

Input: a graph G and an integer k;
Output: all k-vertex connected components;

1: $VCC_k(G) \leftarrow \emptyset$;
2: **while** $\exists u : d(u) < k$ **do** remove u and incident edges;
3: identify connected components $\mathcal{G} = \{G_1, G_2, ..., G_t\}$ in G;
4: **for all** connected component $G_i \in \mathcal{G}$ **do**
5: $\quad\mathcal{S} \leftarrow$ GLOBAL-CUT(G_i, k);
6: $\quad$**if** $\mathcal{S} = \emptyset$ **then**
7: $\quad\quad VCC_k(G) \leftarrow VCC_k(G) \cup \{G_i\}$;
8: $\quad$**else**
9: $\quad\quad\mathcal{G}_i \leftarrow$ OVERLAP-PARTITION$(G_i, \mathcal{S})$;
10: $\quad\quad$**for all** $G_i^j \in \mathcal{G}_i$ **do**
11: $\quad\quad\quad VCC_k(G) \leftarrow VCC_k(G) \cup$ KVCC-ENUM(G_i^j, k);
12: **return** $VCC_k(G)$;

13: **Procedure** OVERLAP-PARTITION(Graph G', Vertex Cut $\mathcal{S}$)
14: $\mathcal{G} \leftarrow \emptyset$;
15: remove vertices in $\mathcal{S}$ and their adjacent edges from G';
16: **for all** connected component G_i' of G' **do**
17: $\quad\mathcal{G} \leftarrow \mathcal{G} \cup \{G'[V(G_i') \cup \mathcal{S}]\}$;
18: **return** $\mathcal{G}$;

3. ALGORITHM FRAMEWORK

3.1 The Cut-Based Framework

To compute all k-VCCs in a graph, we introduce a cut-based framework KVCC-ENUM in this section. We define *vertex cut*.

DEFINITION 4. (VERTEX CUT) *Given a connected graph G, a vertex subset $\mathcal{S} \subset V$ is a vertex cut if the removal of $\mathcal{S}$ from G results in a disconnected graph.*

From Definiton 4, we know that the vertex cut may not be unique for a given graph G, and the vertex connectivity is the size of the minimum vertex cut. For a complete graph, there is no vertex cut since any two vertices are adjacent. The size of a vertex cut is the number of vertices in the cut. In the rest of paper, we use cut to represent vertex cut when the context is clear.

The Algorithm. Given a graph G, the general idea of our cut-based framework is given as follows. If G is k-connected, G itself is a k-VCC. Otherwise, there must exist a qualified cut $\mathcal{S}$ whose size is less than k. In this case, we find such cut and partition G into overlapped subgraphs using the cut. We repeat the partition procedure until each remaining subgraph is a k-VCC. From Theorem 3, we know that a k-VCC must be a k-core (a graph with minimum degree no smaller than k). Thus we can compute all k-cores in advance to reduce the size of the graph.

The pseudocode of our framework is presented in Algorithm 1. In line 2, the algorithm computes the k-core by iteratively removing the vertices whose degree is less than k and terminates once no such vertex exists. Then we identify connected components of the input graph G. For each connected component G_i (line 4), we first find a cut of G_i by invoking the subroutine GLOBAL-CUT (line 5). Here, we only need to find a cut with fewer than k vertices instead of a minimum cut. The detailed implementation of GLOBAL-CUT will be introduced later. If there is no such cut, it means G_i is k-connected and we add it to the result list $VCC_k(G)$ (line 6-7). Otherwise, we partition the graph into overlapped subgraphs using the cut $\mathcal{S}$ by invoking OVERLAP-PARTITION (line 9). We recursively cut each of other subgraphs using the same procedure KVCC-ENUM (line 11) until all remaining subgraphs are k-VCCs. Next, we introduce the subroutine OVERLAP_PARTITION, which partitions the graph into overlapped subgraphs by cut $\mathcal{S}$.

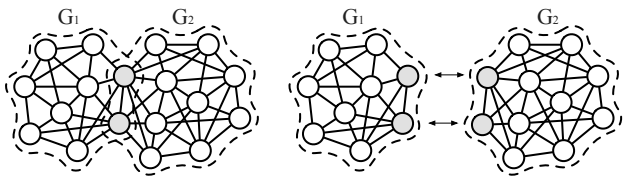

Figure 2: An example of overlapped graph partition.

Overlapped Graph Partition. To partition a graph G into overlapped subgraphs using a cut $\mathcal{S}$, we cannot simply remove all vertices in $\mathcal{S}$, since such vertices may be the overlapped vertices of two or more k-VCCs. Subroutine OVERLAP-PARTITION is shown in line 13-18 of Algorithm 1. We first remove the vertices in $\mathcal{S}$ along with their adjacent edges from G'. G' will become disconnected after removing $\mathcal{S}$, since $\mathcal{S}$ is a vertex cut of G'. We can simply add the cut $\mathcal{S}$ into each connected component G_i' of G' and return induced subgraph $G'[V(G_i') \cup \mathcal{S}]$ as the partitioned subgraph (line 17-18). Partitioned subgraphs overlap each other since the cut $\mathcal{S}$ is duplicated in these subgraphs. Below, we use an example to illustrate the partition process.

EXAMPLE 3. *We consider a graph G on the left of Fig. 2. given the input parameter $k = 3$, we can find a vertex cut in which all vertices are marked by gray. These vertices belong to both 3-VCCs, G_1 and G_2. Thus, given a cut $\mathcal{S}$ of graph G, we partition the graph by duplicating the induced subgraph of cut $\mathcal{S}$. As shown on the right of Fig. 2, we obtain two 3-VCCs, G_1 and G_2, by duplicating the two cut vertices and their inner edges.*

3.2 Algorithm Correctness

In this section, we prove the correctness of KVCC-ENUM using the following lemmas.

LEMMA 1. *Each of the subgraphs returned by* KVCC-ENUM *is k-vertex connected.*

PROOF. We prove it by contradiction. Assume one of the result subgraphs G_i is not k-connected. GLOBAL-CUT in line 5 will find a vertex cut. G_i will be partitioned in line 9 and cannot be returned, which contradicts that G_i is in the result list. □

LEMMA 2. *(Completeness) The result returned by* KVCC-ENUM *contains all k-VCCs of the input graph G.*

PROOF. Suppose graph G is partitioned into overlapped subgraphs $\mathcal{G}' = \{ G_1', G_2', \ldots\}$ using a vertex cut $\mathcal{S}$. We first prove that each k-VCC G_i of G is contained in at least one subgraph in $\mathcal{G}'$. We prove this by contradiction. We suppose that G_i is not contained in any subgraph in $\mathcal{G}'$. Consider the computation of $\mathcal{G}'$, after we remove the vertices in $\mathcal{S}$ and their adjacent edges from G_i, the remaining vertices in G_i are contained in at least two graphs in $\mathcal{G}'$. This indicates that $\mathcal{S}$ is a vertex cut of G_i. Since $|\mathcal{S}| < k$, G_i cannot be a k-VCC, which contradicts that G_i is a k-VCC. Therefore, we prove that we will not lose any k-VCC. From Lemma 1, we know that each of the returned subgraphs of KVCC-ENUM is k-connected. Therefore, all maximal subgraphs that are k-connected will be returned by KVCC-ENUM. In other words, all k-VCCs will be returned by KVCC-ENUM. □

LEMMA 3. *(Redundancy-Free) There does not exist two subgraphs G_i and G_j returned by* KVCC-ENUM *such that $G_i \subseteq G_j$.*

PROOF. We prove it by contradiction. Suppose there are two subgraphs G_i and G_j returned by KVCC-ENUM such that $G_i \subseteq$

G_j. On the one hand, we have $|V(G_i) \cap V(G_j)| = |V(G_i)| \geq k$. On the other hand, there must exist a partition $\mathcal{G}' = \{ G'_1, G'_2, \ldots\}$ of G by a certain cut $\mathcal{S}$ such that G_i and G_j are contained in two different graphs in $\mathcal{G}'$. From the partition procedure, we know that G_i and G_j have at most $k-1$ common vertices. This contradicts $|V(G_i) \cap V(G_j)| \geq k$. Therefore, the lemma holds. □

THEOREM 4. KVCC-ENUM *correctly computes all k-VCCs of G.*

PROOF. From Lemma 1, we know that all subgraphs returned by KVCC-ENUM are k-connected. From Lemma 2, we know that all k-VCCs are returned by KVCC-ENUM. From Lemma 3, we know that all k-connected subgraphs returned by KVCC-ENUM are maximal, and no redundant subgraph will be produced. Therefore, KVCC-ENUM (Algorithm 1) correctly computes all k-VCCs of G. □

Next, we show how to efficiently compute all k-VCCs following the framework in Algorithm 1. From Algorithm 1, we know that the key to improving algorithmic efficiency is to efficiently compute the vertex cut of a graph G. Below, we first introduce a basic algorithm in Section 4 to compute the vertex cut of a graph in polynomial time, and then we explore optimization strategies to accelerate the computation of the vertex cut in Section 5.

4. BASIC SOLUTION

In the previous section, we propose a cut-based framework named KVCC-ENUM to compute all k-VCCs. A key step in Algorithm 1 is GLOBAL-CUT. Before giving the detailed implementation of GLOBAL-CUT, we discuss techniques to find the edge-cut, which is highly related to the vertex-cut. Here, an edge-cut is a set of edges the removal of which will make the graph disconnected. We will show that these methods cannot be directly used to find the vertex-cut.

Maximum Flow. A basic solution to find edge cut is the maximum flow algorithm. With a given maximum flow, we can easily compute a minimum edge cut based on the *Max-Flow Min-Cut Theorem*. However, the flow algorithm only considers capacity of each edge and does not have any limitation on that of vertex, which is obviously not suitable for finding the vertex cut.

Min Edge-Cut. Stoer and Wagner [25] proposed an algorithm to find global minimum edge cut in an undirected graph. The general idea is iteratively finding an edge-cut and merging a pair of vertices. It returns the edge-cut with the smallest value after $n-1$ merge operations. Given an upper bound k, the algorithm terminates once an edge-cut with fewer then k edges is found. However, this algorithm is not suitable for finding the vertex-cut since we do not know whether a vertex is included in the cut or not. Therefore, we cannot simply merge any two vertices in the whole procedure.

4.1 Find Vertex Cut

We give some necessary definitions before introducing the idea to implement GLOBAL-CUT.

DEFINITION 5. (MINIMUM u-v CUT) *A vertex cut $\mathcal{S}$ is a u-v cut if u and v are in disjoint subsets after removing $\mathcal{S}$, and it is a minimum u-v cut if its size is no larger than that of other u-v cuts.*

DEFINITION 6. (LOCAL CONNECTIVITY) *Given a graph G, the local connectivity of two vertices u and v, denoted by $\kappa(u, v, G)$, is defined as the size of the minimum u-v cut. $\kappa(u, v, G) = +\infty$ if no such cut exists.*

Algorithm 2 GLOBAL-CUT(G, k)

Input: a graph G and an integer k;
Output: a vertex cut with fewer than k vertices;

1: compute a sparse certification $\mathcal{SC}$ of G;
2: select a source vertex u with minimum degree;
3: construct the directed flow graph $\overline{\mathcal{SC}}$ of $\mathcal{SC}$;
4: **for all** $v \in V$ **do**
5: $\quad \mathcal{S} \leftarrow$ LOC-CUT$(u, v, \overline{\mathcal{SC}}, \mathcal{SC})$;
6: $\quad$ **if** $\mathcal{S} \neq \emptyset$ **then return** $\mathcal{S}$;
7: **for all** $v_a \in N(u)$ **do**
8: $\quad$ **for all** $v_b \in N(u)$ **do**
9: $\quad\quad \mathcal{S} \leftarrow$ LOC-CUT$(v_a, v_b, \overline{\mathcal{SC}}, \mathcal{SC})$;
10: $\quad\quad$ **if** $\mathcal{S} \neq \emptyset$ **then return** $\mathcal{S}$;
11: return $\emptyset$;

12: **Procedure** LOC-CUT$(u, v, \overline{G}, G)$
13: **if** $v \in N(u)$ or $v = u$ **then return** $\emptyset$;
14: $\lambda \leftarrow$ calculate the maximum flow from u to v in $\overline{G}$;
15: **if** $\lambda \geq k$ **then return** $\emptyset$;
16: compute the minimum edge cut in $\overline{G}$;
17: **return** the corresponding vertex cut in G;

Based on Definiton 6, we define two local k connectivity relations as follows:

- $u \equiv^k_G v$: The local connectivity between u and v is not less than k in graph G, i.e., $\kappa(u, v, G) \geq k$.
- $u \not\equiv^k_G v$: The local connectivity between u and v is less than k in graph G, i.e., $\kappa(u, v, G) < k$.

We omit the suffix G, and use $u \equiv^k v$ and $u \not\equiv^k v$ to denote $u \equiv^k_G v$ and $u \not\equiv^k_G v$ respectively when the context is clear. Once $u \equiv^k v$, we say u and v is k-local connected. Obviously, $u \equiv^k v$ and $v \equiv^k u$ are equivalent.

The GLOBAL-CUT Algorithm. We follow [12] to implement GLOBAL-CUT. Given a graph G, we assume that G contains a vertex cut $\mathcal{S}$ such that $|\mathcal{S}| < k$. We consider an arbitrary source vertex u. There are only two cases: (i) $u \notin \mathcal{S}$ and (ii) $u \in \mathcal{S}$. The general idea of algorithm GLOBAL-CUT considers two cases. In the first phase, we select a vertex u and test the local connectivity between u and all other vertices v in G. We have either (a) $u \in \mathcal{S}$ or (b) G is k-connected if each local connectivity is not less than k. In the second phase, we consider the case $u \in \mathcal{S}$ and test the local connectivity between any two neighbors of u based on Lemma 4. More details can be found in [12].

LEMMA 4. *Given a non-k-vertex connected graph G and a vertex $u \in \mathcal{S}$ where $\mathcal{S}$ is a vertex cut and $|\mathcal{S}| < k$, there exist $v, v' \in N(u)$ such that $v \not\equiv^k v'$.*

The pseudocode is given in Algorithm 2. An optimization here is computing a sparse certificate of the original graph in line 1. Given a graph $G(V, E)$, a sparse certificate is a subset of edges $E' \in E$, such that the subgraph $G'(V, E')$ is k-connected if and only if G is k-connected. Undoubtedly, the same algorithm is more efficient in a sparser graph. We will introduce the details of sparse certification in the next subsection.

The first phase is shown in line 4-6. Once finding such cut $\mathcal{S}$, we return it as the result. Similarly, the second phase is shown in line 7-10. Here, the procedure LOC-CUT tests the local connectivity between u and v and returns the vertex cut if $u \not\equiv^k v$ (line 5 and line 9). To invoke LOC-CUT, we need to transform the original graph G into a directed flow graph $\overline{G}$. The details on how to construct the directed flow graph are introduced as follows.

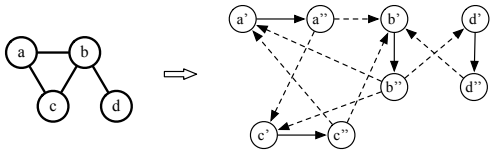

Figure 3: An example of directed flow graph construction.

Directed Flow Graph. The directed flow graph $\overline{G}$ of graph G is an auxiliary directed graph which is used to calculate the local connectivity between two vertices. Given a graph G, we can construct the directed flow graph as follows. Each vertex u in G is represented by an directed edge e_u in the directed flow graph $\overline{G}$. Let u' and u'' denote the starting vertex and ending vertex of e_u. For each edge (u, v) in G, we construct two directed edges: One is from u'' to v', and the other is from v'' to u'. Consequently, we obtain $\overline{G}$ with $2n$ vertices and $n + 2m$ edges and the capacity of every edge is 1.

EXAMPLE 4. *Fig. 3 gives an example of the directed flow graph construction. The solid lines in the directed flow graph represent vertices in the original graph, and the dashed lines in the directed flow graph represent edges in the original graph. The original graph contains 4 vertices and 4 edges, and the directed flow graph contains 8 vertices and 12 edges.*

The LOC-CUT Procedure. By using the directed flow graph, we convert vertex connectivity problem into edge connectivity problem. To calculate the local connectivity of two vertices u and v, we perform the maximum flow algorithm on the directed flow graph. The value of the maximum flow is the local connectivity between u and v.

The pseudocode of LOC-CUT is given form line 12 to line 17 in Algorithm 2. It first checks whether v is a neighbor of u in line 13. If $u \in N(v)$, we always have $u \equiv^k v$ because of Lemma 5.

LEMMA 5. *$u \equiv^k v$ if $(u, v) \in E$.*

Then the procedure computes the maximum flow λ from u to v in $\overline{G}$ in line 14. If $\lambda \geq k$, we have $u \equiv^k v$ and the procedure returns $\emptyset$ in line 15. Otherwise, we compute the edge cut in $\overline{G}$ in line 16. Then we locate the corresponding vertices in the original graph G for each edge in the edge cut and return them as the vertex cut of G (line 16-17).

4.2 Sparse Certificate

We introduce the details of sparse certificate [8] in this section. In Section 5, we will show that the sparse certificate can not only be used to reduce the graph size, but also used to further reduce the local connectivity testings.

DEFINITION 7. (CERTIFICATE) *A certificate for the k-vertex connectivity of G is a subset E' of E such that the subgraph (V, E') is k-vertex connected if and only if G is k-vertex connected.*

DEFINITION 8. (SPARSE CERTIFICATE) *A certificate for k-vertex connectivity of G is called sparse if it has $O(k \cdot n)$ edges.*

From the definitions, we can see that a sparse certificate is equivalent to the original graph w.r.t k-vertex connectivity. It can also bound the edge size. We compute the sparse certificate (line 1 of Algorithm 2) according to the following theorem.

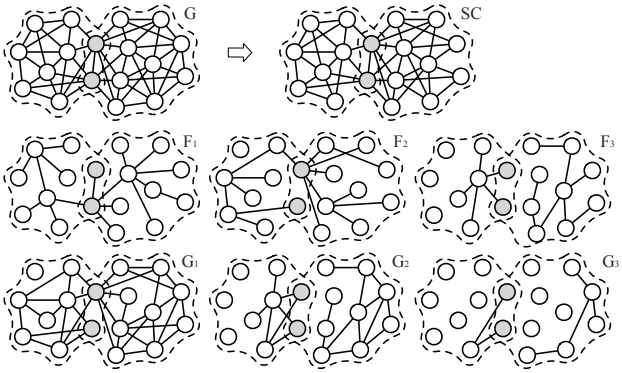

Figure 4: The sparse certificate of given graph G with $k = 3$

THEOREM 5. *Let $G(V, E)$ be an undirected graph and let n denote the number of vertices. Let k be a positive integer. For $i = 1, 2, ..., k$, let E_i be the edge set of a scan first search forest F_i in the graph $G_{i-1} = (V, E - (E_1 \cup E_2 \cup ... \cup E_{i-1}))$. Then $E_1 \cup E_2 \cup ... \cup E_k$ is a certificate for the k-vertex connectivity of G, and this certificate has at most $k \times (n - 1)$ edges [8].*

Based on Theorem 5, we can simply generate the sparse certificate of G using scan first search k times, each of which creates a scan first search forest F_i. Below, we introduce how to perform a scan first search.

Scan First Search. In a scan first search of given graph G, for each connected component, we start from scanning a root vertex by marking all its neighbors. We scan an arbitrary marked but unscanned vertex each time and mark all its unvisited neighbors. This step is performed until all vertices are scanned. The resulting search forest forms the scan first forest of G. Obviously, a breath first search is a special case of scan first search.

EXAMPLE 5. *Fig. 4 presents construction of a sparse certificate for the graph G. Let $k = 3$. For $i \in \{1, 2, 3\}$, F_i denotes the scan first search forest obtained from G_{i-1}. G_i is obtained by removing the edges in F_i from G_{i-1}. G_0 is the input graph G. The obtained sparse certificate SC is shown on the right side of G with $SC = F_1 \cup F_2 \cup F_3$. All removed edges are shown in G_3.*

4.3 Algorithm Analysis

We analyze the basic algorithm in this section. In the directed flow graph, all edge capacities are equal to 1 and every vertex either has a single edge emanating from it or has a single edge entering it. For this kind of graph, the time complexity for computing the maximum flow is $O(n^{1/2}m)$ [14]. Note that we do not need to calculate the exact flow value in the algorithm. Once the flow value reaches k, we know that local connectivity between any two given vertices is at least k and we can terminate the maximum flow algorithm. The time complexity for the flow computation is $O(\min(n^{1/2}, k) \cdot m)$. Given a flow value and corresponding residual network, we can perform a depth first search to find the cut. It costs $O(m + n)$ time. As a result, we have the following lemma:

LEMMA 6. *The time complexity of* LOC-CUT *is* $O(\min(n^{1/2}, k) \cdot m)$.

Next we discuss the time complexity of GLOBAL-CUT. The construction of both sparse certificate and directed flow graph costs $O(m+n)$ CPU time. Let δ denote the minimum degree in the input graph. We can easily get following lemma.

LEMMA 7. GLOBAL-CUT *invokes* LOC-CUT $O(n + \delta^2)$ *times in the worst case.*

Next we discuss the CPU time complexity of the entire algorithm KVCC-ENUM. KVCC-ENUM iteratively removes vertices with degree less than k in line 2. This costs $O(m+n)$ time. Identifying all connected components can be performed by adopting a depth first search (line 3). This also need $O(m + n)$ time. To study the total time complexity spent by invoking GLOBAL-CUT, we first give the following lemma.

LEMMA 8. *For each subgraph created by the overlapped partition, at most* $k - 1$ *vertices and* $\frac{(k-1)(k-2)}{2}$ *edges are increased after the partition.*

PROOF. The vertex cut $\mathcal{S}$ contains not more than $k-1$ vertices, and only these vertices exist in the overlapped part. Therefore, at most $\frac{(k-1)(k-2)}{2}$ incident edges are duplicated. □

LEMMA 9. *Given a graph* G *and an integer* k*, for each connected component* C *obtained by overlapped partition in Algorithm 1,* $|V(C)| \geq k + 1$.

PROOF. Let $\mathcal{S}$ denote a vertex cut in an overlapped partition. C is one of the connected components obtained in this partition. Let H denote the vertex set of all vertices in $V(C)$ but not in $\mathcal{S}$, i.e., $H = \{u | u \in V(C), u \notin \mathcal{S}\}$. We have $H \neq \emptyset$. Note that each vertex in the graph has a degree at least k in G (line 5 in Algorithm 1). There exist at least k neighbors for each vertex u in H and therefore for each neighbor v of u we have $v \in C$ according to Lemma 5. Thus, we have $|V(C)| \geq k + 1$. □

LEMMA 10. *Given a graph* G *and an integer* k*, the total number of overlapped partitions during the algorithm* KVCC-ENUM *is no larger than* $\frac{n-k-1}{2}$.

PROOF. Suppose that λ is the total number of overlapped partitions during the whole algorithm KVCC-ENUM. This generates at least $\lambda + 1$ connected components. We know from Lemma 9 that each connected component contains at least $k + 1$ vertices. Thus, we have at least $(\lambda + 1)(k + 1)$ vertices in total.

On the other hand, we increase at most $k-1$ vertices in each subgraph obtained by an overlapped partition according to Lemma 8. Thus, at most $\lambda(k-1)$ vertices are added. We obtain the following formula.

$$(\lambda + 1)(k + 1) \leq n + \lambda(k - 1)$$

Rearranging the formula, we have $\lambda \leq \frac{n-k-1}{2}$. □

Next, we prove the upper bound for number of k-VCCs.

THEOREM 6. *Given a graph* G *and an integer* k*, there are at most* $\frac{n}{2}$ *k-VCCs, i.e.,* $|VCC_k(G)| < \frac{|V(G)|}{2}$.

PROOF. Similar to the proof of Lemma 10, let λ be the times of overlapped partitions in the whole algorithm KVCC-ENUM. At most $\lambda(k-1)$ vertices are increased. Let σ be the number of connected components obtained in all partitions. We have $\sigma > \lambda$. Each connected component contains at least $k + 1$ vertices according to Lemma 9. Note that each connected component is either a k-VCC or a graph that does not contain any k-VCC. Otherwise, the connected component will be further partitioned. Let x be the number of k-VCCs and y be the number of connected components that do not contain any k-VCC, i.e., $x + y = \sigma$. We know that a k-VCC contains at least $k + 1$ vertices. Thus there are at least $x(k+1) + y(k+1)$ vertices after finishing all partitions. We have following formula.

$$x(k + 1) + y(k + 1) \leq n + \lambda(k - 1)$$

Since $\lambda < \sigma$ and $\sigma = x+y$, we rearrange the formula as follows.

$$x(k + 1) + y(k + 1) < n + x(k - 1) + y(k - 1)$$
$$x(k + 1) < n + x(k - 1)$$

Therefore, we have $x < \frac{n}{2}$. □

THEOREM 7. *The total time complexity of* KVCC-ENUM *is* $O(\min(n^{1/2}, k) \cdot m \cdot (n + \delta^2) \cdot n)$.

PROOF. The total time complexity of KVCC-ENUM is dependent on the number of times GLOBAL-CUT is invoked. Suppose GLOBAL-CUT is invoked p times during the whole KVCC-ENUM algorithm, the number of overlapped partitions during the whole KVCC-ENUM algorithm is p_1 and the total number of k-VCCs is p_2. It is easy to see that $p = p_1 + p_2$. From Lemma 10, we know that $p_1 \leq \frac{n-k-1}{2} < \frac{n}{2}$. From Theorem 6, we know that $p_2 < \frac{n}{2}$. Therefore, we have $p = p_1 + p_2 < n$. According to Lemma 6 and Lemma 7, the total time complexity of KVCC-ENUM is $O(\min(n^{1/2}, k) \cdot m \cdot (n + \delta^2) \cdot n)$. □

Discussion. Theorem 7 shows that all k-VCCs can be enumerated in polynomial time. Although the time complexity is still high, it performs much better in practice. Note that the time complexity is the product of three parts:

- The first part $O(\min(n^{1/2}, k) \cdot m)$ is the time complexity for LOC-CUT to test whether there exists a vertex cut of size smaller than k. In practice, the graph to be tested is much smaller than the original graph G since (1) The graph to be tested has been pruned using the k-core technique and sparse certification technique. (2) Due to the graph partition scheme, the input graph is partitioned into many smaller graphs.
- The second part $O(n + \delta^2)$ is the number of times such that LOC-CUT (local connectivity testing) is invoked by the algorithm GLOBAL-CUT. We will discuss how to significantly reduce the number of local connectivity testings in Section 5.
- The third part $O(n)$ is the number of times GLOBAL-CUT is invoked. In practice, the number can be significantly reduced since the number of k-VCCs is usually much smaller than $\frac{n}{2}$.

In the next section, we will explore several search reduction techniques to speed up the algorithm.

5. SEARCH REDUCTION

In the previous section, we introduce our basic algorithm. Recall that in the worst case, we need to test local connectivity between the source vertex u and all other vertices in G using LOC-CUT inGLOBAL-CUT, and we also need to test local connectivity for every pair of neighbors of u. For each pair of vertices, we need to compute the maximum flow in the directed flow graph. Therefore, the key to improving the algorithm is to reduce the number of local connectivity testings (LOC-CUT). In this section, we propose several techniques to avoid unnecessary testings. We can avoid testing local connectivity of a vertex pair (u, v) if we can guarantee that $u \equiv^k v$. We call such operation a sweep operation. Below, we introduce two ways to efficiently prune unnecessary testings, namely neighbor sweep and group sweep, in Section 5.1 and Section 5.2 respectively.

5.1 Neighbor Sweep

In this section, we propose a neighbor sweep strategy to prune unnecessary local connectivity testings (LOC-CUT) in the first phase

of GLOBAL-CUT. Generally speaking, given a source vertex u, for any vertex v, we aim to skip testing the local connectivity of (u, v) according to the information of the neighbors of v. Below, we explore two neighbor sweep strategies, namely neighbor sweep using side-vertex and neighbor sweep using vertex deposit.

5.1.1 Neighbor Sweep using Side-Vertex

We first define *side-vertex* as follows.

DEFINITION 9. (SIDE-VERTEX) *Given a graph G and an integer k, a vertex u is called a side-vertex if there does not exist a vertex cut $\mathcal{S}$ such that $|\mathcal{S}| < k$ and $u \in \mathcal{S}$.*

Based on Definiton 9, we give the following lemma to show the transitive property regarding the local k connectivity relation $\equiv^k$.

LEMMA 11. *Given a graph G and an integer k, suppose $a \equiv^k b$ and $b \equiv^k c$, we have $a \equiv^k c$ if b is a side-vertex.*

PROOF. We prove it by contradiction. Assume that b is a side-vertex and $a \not\equiv^k c$. There exists a vertex cut with $k-1$ or fewer vertices between a and c. b is not in any such cut since it is a side-vertex. Then we have either $b \not\equiv^k a$ or $b \not\equiv^k c$. This contradicts the precondition that $a \equiv^k b$ and $b \equiv^k c$. □

A wise way to use the transitive property of the local connectivity relation in Lemma 11 can largely reduce the number of unnecessary testings. Consider a selected source vertex u in algorithm GLOBAL-CUT. We assume that LOC-CUT (line 5) returns $\emptyset$ for a vertex v, i.e., $u \equiv^k v$. We know from Lemma 11 that the vertex pair (u, w) can be skipped for local connectivity testing if (i) $v \equiv^k w$ and (ii) v is a side-vertex. For condition (i), we can use a simple necessary condition according to Lemma 5, that is, for any vertices v and w, $v \equiv^k w$ if $(v, w) \in E$. In the following, we focus on condition (ii) and look for necessary conditions to efficiently check whether a vertex is a side-vertex.

Side-Vertex Detection. To check whether a vertex is a side-vertex, we can easily obtain the following lemma based on Definiton 9.

LEMMA 12. *Given a graph G, a vertex u is a side-vertex if and only if $\forall v, v' \in N(u)$, $v \equiv^k v'$.*

Recall that two vertices are k-local connected if they are neighbors of each other. For the k-local connectivity of non-connected vertices, we give another necessary condition below.

LEMMA 13. *Given two vertices u and v, $u \equiv^k v$ if $|N(u) \cap N(v)| \geq k$.*

PROOF. u and v cannot be disjoint after removing any $k-1$ vertices since they have at least k common neighbors. Thus u and v must be k-local connected. □

Combining Lemma 12 and Lemma 13, we derive the following necessary condition to check whether a vertex is a side-vertex.

THEOREM 8. *A vertex u is a side-vertex if $\forall v, v' \in N(u)$, either $(v, v') \in E$ or $|N(v) \cap N(v')| \geq k$.*

PROOF. The theorem can be easily verified using Lemma 5, Lemma 12 and Lemma 13. □

DEFINITION 10. (STRONG SIDE-VERTEX) *A vertex u is called a strong side-vertex if it satisfies the conditions in Theorem 8.*

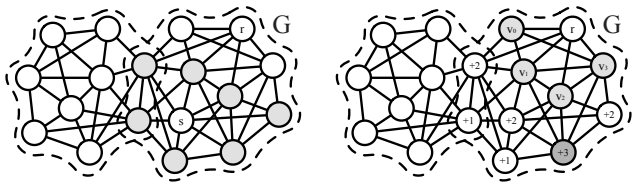

(a) A strong side-vertex s (b) Vertex deposit

Figure 5: Strong side-vertex and vertex deposit when $k = 3$

Using strong side-vertex, we can define our first rule for neighbor sweep as follows.

(Neighbor Sweep Rule 1) *Given a graph G and an integer k, let u be a selected source vertex in algorithm* GLOBAL-CUT *and v be a strong side-vertex in the graph. We can skip the local connectivity testings of all pairs of (u, w) if we have $u \equiv^k v$ and $w \in N(v)$.*

We give an example to demonstrate *neighbor sweep rule 1* below.

EXAMPLE 6. *Fig. 5 (a) presents a strong side-vertex s in graph G while parameter $k = 3$. Assume that r is the source vertex. Any two neighbors of s are either connected by an edge or have at least 3 common neighbors. If first test the local connectivity between r and s and $r \equiv^k s$, we can safely sweep all neighbors of s, which are marked by the gray color in Fig. 5 (a).*

Below, we discuss how to efficiently detect the strong side-vertices and maintain strong side-vertices while the graph is partitioned in the whole algorithm.

Strong Side-Vertex Computation. Following Theorem 8, we can compute all strong side-vertices v in advance and skip all neighbors of v once v is k connected with the source vertex (line 5 in GLOBAL-CUT). We can derive the following lemma.

LEMMA 14. *The time complexity of computing all strong side-vertices in graph G is $O(\sum_{w \in V(G)} d(w)^2)$.*

PROOF. To compute all strong side-vertices in a graph G, we first check all 2-hop neighbors v for each vertex u. Since v and u share a common vertex of 1-hop neighbor, we can easily obtain all vertices which have k common neighbors with u. Any vertex w is considered as 1-hop neighbor of other vertices u $d(w)$ times. We use $d(w)$ steps to obtain 2-hop neighbors of u which share a common vertex w with u. This phase costs $O(\sum_{w \in V(G)} d(w)^2)$ time.

Now for each given vertex u, we have all vertices v sharing k common neighbors with it. For each vertex w, we check whether any two neighbors of w have k common neighbors. This phase also costs $O(\sum_{w \in V(G)} d(w)^2)$ time. Consequently, the total time complexity is $O(\sum_{w \in V(G)} d(w)^2)$. □

After computing all strong side-vertices for the original graph G, we do not need to recompute the strong side-vertices for all vertices in the partitioned graph from scratch. Instead, we can find possible ways to reduce the number of strong side-vertex checks by making use of the already computed strong side-vertices in G. We can do this based on Lemma 15 and Lemma 16 which are used to efficiently detect non-strong side-vertices and strong side-vertices respectively.

LEMMA 15. *Let G be a graph and G_i be one of the graphs obtained by partitioning G using* OVERLAP-PARTITION *in Algorithm 1, a vertex is a strong side-vertex in G if it is a strong side-vertex in G_i.*

PROOF. The strong side-vertex u requires at least k common neighbors between any two neighbors of u. The lemma is obvious since G contains all edges and vertices in G_i. □

From Lemma 15, we know that a vertex is not a strong side-vertex in G_i if it is not a strong side-vertex in G. This property allows us checking limited number of vertices in G_i, which is the set of strong side-vertices in G.

LEMMA 16. *Let G be a graph, G_i be one of the graphs obtained by partitioning G using* OVERLAP-PARTITION *in Algorithm 1, and $\mathcal{S}$ is a vertex cut of G, for any vertex $v \in V(G_i)$, if v is a strong side-vertex in G and $N(v) \cap \mathcal{S} = \emptyset$, then v is also a strong side-vertex in G_i.*

PROOF. The qualification of a strong side-vertex of vertex v requires the information about two-hop neighbors of v. Vertices in $\mathcal{S}$ are duplicated when partitioning the graph. Given a strong side-vertex v in G, if $N(v) \cap \mathcal{S} = \emptyset$, the two-hop neighbors of v are not affected by the partition operation, thus the relationships between the vertices in $N(v)$ are not affected by the partition operation. Therefore, v is still a strong side-vertex in G_i according to Definiton 10. □

With Lemma 15 and Lemma 16, in a graph G_i partitioned from graph G by vertex cut $\mathcal{S}$, we can reduce the scope of strong side-vertex checks from the vertices in the whole graph G_i to the vertices u satisfying following two conditions simultaneously:

- u is a strong side-vertex in G; and
- $N(u) \cap \mathcal{S} \neq \emptyset$.

5.1.2 Neighbor Sweep using Vertex Deposit

Vertex Deposit. The strong side-vertex strategy heavily relies on the number of strong side-vertices. Next, we investigate a new strategy called vertex deposit, to further sweep vertices based on neighbor information. We first give the following lemma:

LEMMA 17. *Given a source vertex u in graph G, for any vertex $v \in V(G)$, we have $u \equiv^k v$ if there exist k vertices $w_1, w_2, \ldots, w_k$ such that $u \equiv^k w_i$ and $w_i \in N(v)$ for any $1 \leq i \leq k$.*

PROOF. We prove it by contradiction. Assume that $u \not\equiv^k v$. There exists a vertex cut $\mathcal{S}$ with $k-1$ or fewer vertices between u and v. For any $w_i (1 \leq i \leq k)$, we have $w_i \equiv^k v$ since $w_i \in N(v)$ (Lemma 5) and we also have $w_i \equiv^k u$. Since $u \not\equiv^k v$, w_i cannot satisfy both $w_i \equiv^k u$ and $w_i \equiv^k v$ unless $w_i \in \mathcal{S}$. Therefore, we obtain a cut $\mathcal{S}$ with at least k vertices w_1, w_2, ..., w_k. This contradicts $|\mathcal{S}| < k$. □

Based on Lemma 17, given a source vertex u, once we find a vertex v with at least k neighbors w_i with $u \equiv^k w_i$, we can obtain $u \equiv^k v$ without testing the local connectivity of (u, v). To efficiently detect such vertices v, we define the deposit of a vertex v as follows.

DEFINITION 11. *(Vertex Deposit) Given a source vertex u, the deposit for each vertex v, denoted by $deposit(v)$, is the number of neighbors w of v such that the local connectivity of w and u has been computed with $w \equiv^k u$.*

According to Definiton 11, suppose u is the source vertex and for each vertex v, $deposit(v)$ is a dynamic value depending on the number of processed vertex pairs. To maintain the vertex deposit, we initialize the deposit to 0 and once we know $w \equiv^k u$ for a certain vertex w, we can increase the deposit $deposit(v)$ for each vertex $v \in N(w)$ by 1. We can obtain the following theorem according to Lemma 17.

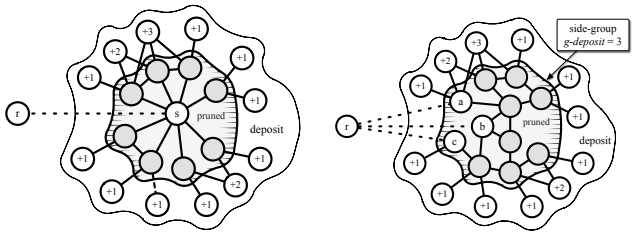

(a) 2-hop deposit (b) group sweep and deposit

Figure 6: Increasing deposit with neighbor and group sweep

THEOREM 9. *Given a source vertex u, for any vertex v, we have $u \equiv^k v$ if $deposit(v) \geq k$.*

Based on Theorem 9, we can derive our second rule for neighbor sweep as follows.

(Neighbor Sweep Rule 2) *Given a selected source vertex u, we can skip the local connectivity testing of pair (u, v) if $deposit(v) \geq k$.*

We show an example below.

EXAMPLE 7. *Fig. 5 (b) gives an example of our vertex deposit strategy. Given the graph G and parameter $k = 3$, let vertex r be the selected source vertex. We assume that v_0, v_1, v_2 and v_3 are tested vertices. All these vertices are local k-connected with vertex r, i.e., $r \equiv^k v_i, i \in 0, 1, 2, 3$, since v_0, v_1, v_2 and v_3 are neighbors of r. We deposit once for the neighbors of each tested vertex. The deposit value for all influenced vertices are given in the figure. We mark the vertices with deposit no less than 3 by dark gray. The local connectivity testing between r and such a vertex can be skipped.*

To increase the deposit of a vertex v, we only need any neighbor of v is local k-connected with the source vertex u. We can also use vertex deposit strategy when processing the strong side-vertex. Given a source vertex u and a strong side-vertex v, we sweep all $w \in N(v)$ if $u \equiv^k v$ according to the side-vertex strategy. Next we increase the deposit for each non-swept vertex $w' \in N(w)$. In other words, for a strong side-vertex, we can possibly sweep its 2-hop neighbors by combining the two neighbor sweep strategies. An example is given below.

EXAMPLE 8. *Fig. 6 (a) shows the process for a strong side-vertex s. Given a source vertex r, assume s is a strong side-vertex and $r \equiv^k s$. All neighbors of s are swept and all 2-hop neighbors of s increase their deposits accordingly. The increased value of the deposit for each vertex depends on the number of connected vertices that are swept.*

5.2 Group Sweep

The neighbor sweep strategy can only prune unnecessary local connectivity testings in the first phase of GLOBAL-CUT by using the neighborhood information. In this subsection, we introduce a new pruning strategy, namely group sweep, which can prune unnecessary local connectivity testings in a batch manner. In group sweep, we do not limit the skipped vertices to the neighbors of certain vertices. More specifically, we aim to partition vertices into vertex groups and sweep a whole group when it satisfies certain conditions. In addition, our group sweep strategy can also be applied to reduce the unnecessary local connectivity testings in both phases of GLOBAL-CUT.

First, we define a new relation regarding a vertex u and a set of vertices C as follows.

$u \equiv^k C$: For all vertices $v \in C, u \equiv^k v$.

Given a source vertex u and a side-vertex v, we assume $u \equiv^k v$. According to the transitive relation in Lemma 11, we can skip testing the pairs of vertices u and w for all w with $w \equiv^k v$. In our neighbor sweep strategy, we select all neighbors of v as such vertices w, i.e., $u \equiv^k N(v)$. To sweep more vertices each time, we define the side-group.

DEFINITION 12. (SIDE-GROUP) *Given a graph G and an integer k, a vertex set $\mathcal{CC}$ in G is a side-group if $\forall u, v \in C, u \equiv^k v$.*

Note that it is possible that a side-group contains vertices in a certain vertex cut $\mathcal{S}$ with $|\mathcal{S}| < k$. Next, we introduce how to construct the side-groups in graph G, and then discuss our group sweep rules.

Side-Group Construction. Section 4.2 introduces sparse certificate to bound the graph size. Let F_i and G_i be the notations defined in Theorem 5. Assume that G is not k-connected and there exists a vertex cut $\mathcal{S}$ such that $|\mathcal{S}| < k$. According to [8], we have the following lemma.

LEMMA 18. *F_k does not contain a simple tree path P_k whose two end points are in different connected components of $G - \mathcal{S}$.*

Based on Lemma 18, we can obtain the following theorem.

THEOREM 10. *Let $\mathcal{CC}$ denote the vertex set of any connected component in F_k. $\mathcal{CC}$ is a side-group.*

PROOF. Assume that $u \not\equiv^k v$ in $\mathcal{CC}$. All simple paths from u to v will cross the vertex cut $\mathcal{S}$. This contradicts Lemma 18. □

EXAMPLE 9. *Review the construction of a sparse certificate in Fig. 4. Given $k = 3$, two connected components with more than one vertex are obtained in F_3. The number of vertices in the two connected components are 6 and 9 respectively. Each of them is a side-group and any two vertices in the same connected component is local 3-connected. Note that the connected component with 6 vertices contains two vertices in the vertex cut as marked by gray.*

We denote all the side-groups as $\mathcal{CS} = \{\mathcal{CC}_1, \mathcal{CC}_2, \ldots, \mathcal{CC}_t\}$. According to Theorem 10, $\mathcal{CS}$ can be easily computed as a by-product of the sparse certificate. With $\mathcal{CS}$, according to the transitive relation in Lemma 11, we can easily obtain the following pruning rule.

(Group Sweep Rule 1) *Let u be the source vertex in the algorithm* GLOBAL-CUT*, given a side-group $\mathcal{CC}$, if there exists a strong side-vertex $v \in \mathcal{CC}$ such that $u \equiv^k v$, we can skip the local connectivity testings of vertex pairs (u, w) for all $w \in \mathcal{CC} - \{v\}$.*

The above group sweep rule relies on the successful detection of a strong side-vertex in a certain side-group. In the following, we further introduce a deposit based scheme to handle the scenario that no strong side-vertex exists in a certain side-group.

Group Deposit. Similar with the vertex deposit strategy, the group deposit strategy aims to deposit the values in a group level. To show our group deposit scheme, we first introduce the following lemma.

LEMMA 19. *Given a source vertex u, an integer k, and a side-group $\mathcal{CC}$, we have $u \equiv^k \mathcal{CC}$ if $|\{v | v \in \mathcal{C}, u \equiv^k v\}| \geq k$.*

PROOF. We prove it by contradiction. Assume that there exists a vertex w in $\mathcal{CC}$ such that $u \not\equiv^k w$. A vertex cut $\mathcal{S}$ exists with $|\mathcal{S}| < k$. Let $v_0, v_1, ..., v_{k-1}$ be the k vertices in $\mathcal{CC}$ such that $u \equiv^k v_i, 0 \leq i \leq k-1$. We have $w \equiv^k v_i$ based on the definition of a side-group. Each v_i must belong to $\mathcal{S}$ since $u \not\equiv^k w$. As a result, the size of $\mathcal{S}$ is at least k. This contradicts $|\mathcal{S}| < k$. □

Based on Lemma 19, given a source vertex u, once we find a side-group $\mathcal{CC}$ with at least k vertices v with $u \equiv^k v$, we can get $u \equiv^k \mathcal{CC}$ without testing the local connectivity from u to other vertices in $\mathcal{CC}$. To efficiently detect such side-groups $\mathcal{CC}$, we define the group deposit of a side-group $\mathcal{CC}$ as follows.

DEFINITION 13. *(Group Deposit) The group deposit for each side-group $\mathcal{CC}$, denoted by $g\text{-}deposit(\mathcal{CC})$, is the number of vertices $v \in \mathcal{CC}$ such that the local connectivity of v and u has been computed with $v \equiv^k u$.*

According to Definiton 13, suppose u is the source vertex, for each side-group $\mathcal{CC} \in \mathcal{CS}$, $g\text{-}deposit(\mathcal{CC})$ is a dynamic value depending on the already processed vertex pairs. To maintain the group deposit for each side-group $\mathcal{CC}$, we initialize the group deposit for $\mathcal{CC}$ to 0. Once $v \equiv^k u$ for a certain vertex $v \in \mathcal{CC}$, we can increase $g\text{-}deposit(\mathcal{CC})$ by 1. We obtain the following theorem according to Lemma 19.

THEOREM 11. *Given a source vertex u, for any side-group $\mathcal{CC} \in \mathcal{CS}$, we have $u \equiv^k \mathcal{CC}$ if $g\text{-}deposit(\mathcal{CC}) \geq k$.*

Based on Theorem 11, we can derive our second rule for group sweep as follows.

(Group Sweep Rule 2) *Given a selected source vertex u, we can skip the local connectivity testings between u and vertices in $\mathcal{CC}$ if $g\text{-}deposit(\mathcal{CC}) \geq k$.*

Note that a group sweep operation can further trigger a neighbor sweep operation and vice versa, since both operations result in new local k-connected vertex pairs. We show an example below.

EXAMPLE 10. *Fig. 6 (b) presents an example of group sweep. Suppose $k = 3$ and the gray area is a detected side-group. Given a source vertex r, assume that a, b, c are the tested vertices with $r \equiv^k a, r \equiv^k b$ and $r \equiv^k c$ respectively. According to Theorem 11, we can safely sweep all vertices in the same side-group. Also, we apply the vertex deposit strategy for neighbors outside the side-group. The increased value of deposit is shown on each vertex.*

Next we show that the side-groups can also be used to prune the local connectivity testings in the second phase of GLOBAL-CUT. Recall that in the second phase of GLOBAL-CUT, given a source vertex u, we need to test the local connectivity of every pair (v_a, v_b) of the neighbors of u. With side-groups, we can easily obtain the following group sweep rule.

(Group Sweep Rule 3) *Let u be the source vertex, and v_a and v_b be two neighbors of u. If v_a and v_b belong to the same side-group, we have $v_a \equiv^k v_b$ and thus we do not need to test the local connectivity of (v_a, v_b) in the second phase of* GLOBAL-CUT.

The detailed implementation of the neighbor sweep and group sweep techniques is given in the following section.

5.3 The Overall Algorithm

In this section, we combine our pruning strategies and give the implementation of optimized algorithm GLOBAL-CUT*. The pseudocode is presented in Algorithm 3. We can replace GLOBAL-CUT with GLOBAL-CUT* in KVCC-ENUM to obtain our final algorithm to compute all k-VCCs.

The GLOBAL-CUT* algorithm still follows the similar idea of GLOBAL-CUT that consider a source vertex u, and then compute the vertex cut in two phases based on whether u belongs to the vertex cut $\mathcal{S}$. Given a source vertex u, phase 1 (line 8-15) considers the case that $u \notin \mathcal{S}$. Phase 2 (line 16-21) considers the case that

$u \in \mathcal{S}$. If in both phase, the vertex cut $\mathcal{S}$ is not found, there is no such a cut and we simply return $\emptyset$ in line 22.

We compute the side-groups $\mathcal{CS}$ while computing the sparse certificate (line 1). Note that here we only consider the side-group whose size is larger than k, since the group can be swept only if at least k vertices in the group are swept according to Theorem 11. Then we compute all strong side-vertices, $\mathcal{SV}$ based on Theorem 8 (line 3). Here, the strong side-vertices are computed based on the method discussed in Section 5.1.1. If $\mathcal{SV}$ is not empty, we can select one inside vertex as source vertex u and do not need to consider the phase 2, because u cannot be in any cut $\mathcal{S}$ with $|\mathcal{S}| < k$ in this case. Otherwise, we still select the source vertex u with the minimum degree (line 4-7).

In phase 1 (line 8-15), we initialize the group deposit for each side-group, which is number of swept vertices in the side-group, to 0 (line 8). Also, we initialize the local deposit for each vertex to 0 and pru for each vertex to *false* (line 9). Here, pru is used to mark whether a vertex can be swept. Since the source vertex u is local k-connected with itself, we first apply the sweeping rules on the source vertex by invoking SWEEP procedure (line 10). Intuitively, a vertex that is close to the source vertex u tends to be in the same k-VCC with u. In other words, a vertex v that is far away from u tends to be separated from u by a vertex cut $\mathcal{S}$. Therefore, we process vertices v in G according to the non-ascending order of $dist(u, v, G)$ (line 11). We aim to find the vertex cut by processing as few vertices as possible. For each vertex v to be processed in phase 1, we skip it if $pru(v)$ is true (line 12). Otherwise, we test the local connectivity of u and v using LOC-CUT (line 13). If there is a cut $\mathcal{S}$ with size smaller than k, we simply return $\mathcal{S}$ (line 15). Otherwise, we invoke SWEEP procedure to sweep vertices using the sweep rules introduced in Section 5. We will introduce the SWEEP procedure in detail later.

In phase 2 (line 16-21), we first check whether the source vertex u is a strong side-vertex. If so, we can skip phase 2 since a strong side-vertex is not contained in any vertex cut with size smaller than k. Otherwise, we perform pair-wise local connectivity testings for all vertices in $N(u)$. Here, we apply the *group sweep rule 3* and skip testing those pairs of vertices that are in the same side-group (line 19).

Procedure SWEEP**.** The procedure SWEEP is shown in Algorithm 4. To sweep a vertex v, we set $pru(v)$ to be true. This operation may result in neighbor sweep and group sweep of other vertices as follows.

- (Neighbor Sweep) In line 1-5, we consider the neighbor sweep. For all the neighbors w of v that have not been swept, we first increase $deposit(w)$ by 1 based on Definiton 11. Then we consider two cases. The first case is that v is a strong side-vertex. According to *neighbor sweep rule 1* in Section 5.1.1, w can be swept since w is a neighbor of v. The second case is $deposit(w) > k$. According to *neighbor sweep rule 2* in Section 5.1.2, w can be swept. In both cases, we invoke SWEEP to sweep w recursively. (line 4-5)
- (Group Sweep) In line 6-11, we consider the group sweep if v is contained in a side-group $\mathcal{CC}_i$. We first increase $g\text{-}deposit(\mathcal{CC}_i)$ by 1 based on Definiton 13. Then we consider two cases. The first case is that v is a strong side-vertex. According to *group sweep rule 1* in Section 5.2, we can sweep all vertices in $\mathcal{CC}_i$. The second case is that $g\text{-}deposit \geq k$. According to *group sweep rule 2* in Section 5.2, we can sweep all vertices in $\mathcal{CC}_i$. In both cases we recursively invoke SWEEP to sweep each unswept vertex in $\mathcal{CC}_i$ (line 8-11).

Algorithm 3 GLOBAL-CUT$^*(G, k)$

Input: a graph G and an integer k;
Output: a vertex cut with size smaller than k;

1: compute a sparse certification $\mathcal{SC}$ of G and collect all side-groups as $\mathcal{CS} = \{\mathcal{CC}_1, ..., \mathcal{CC}_t\}$;
2: construct the directed flow graph $\overrightarrow{\mathcal{SC}}$ of $\mathcal{SC}$;
3: $\mathcal{SV} \leftarrow$ compute all strong side vertices in $\mathcal{SC}$;
4: **if** $\mathcal{SV} = \emptyset$ **then**
5: select a vertex u with minimum degree;
6: **else**
7: randomly select a vertex u from $\mathcal{SV}$;
8: **for all** $\mathcal{CC}_i$ in $\mathcal{CS}$: $g\text{-}deposit(\mathcal{CC}_i) \leftarrow 0$;
9: **for all** v in V: $deposit(v) \leftarrow 0, pru(v) \leftarrow$ false;
10: SWEEP$(u, pru, deposit, g\text{-}deposit, \mathcal{CS})$;
11: **for all** $v \in V$ in non-ascending order of $dist(u, v, G)$ **do**
12: **if** $pru(v) =$ true **then continue**;
13: $\mathcal{S} \leftarrow$ LOC-CUT$(u, v, \overrightarrow{\mathcal{SC}}, \mathcal{SC})$;
14: **if** $\mathcal{S} \neq \emptyset$ **then return** $\mathcal{S}$;
15: SWEEP$(v, pru, deposit, g\text{-}deposit, \mathcal{CS})$;
16: **if** u is not a strong side-vertex **then**
17: **for all** $v_a \in N(u)$ **do**
18: **for all** $v_b \in N(u)$ **do**
19: **if** v_a and v_b are in the same $\mathcal{CC}_i$ **then continue**;
20: $\mathcal{S} \leftarrow$ LOC-CUT$(u, v, \overrightarrow{\mathcal{SC}}, \mathcal{SC})$;
21: **if** $\mathcal{S} \neq \emptyset$ **then return** $\mathcal{S}$;
22: **return** $\emptyset$;

Algorithm 4 SWEEP$(v, pru, deposit, g\text{-}deposit, \mathcal{CS})$

1: $pru(v) \leftarrow$ true;
2: **for all** $w \in N(v)$ s.t. $pru(w) =$ false **do**
3: $deposit(w)$++;
4: **if** v is a strong side-vertex **or** $deposit(w) \geq k$ **then**
5: SWEEP$(w, pru, deposit, g\text{-}deposit, \mathcal{CS})$;
6: **if** v is contained in a $\mathcal{CC}_i$ and $\mathcal{CC}_i$ has not been processed **then**
7: $g\text{-}deposit(\mathcal{CC}_i)$++;
8: **if** v is a strong side-vertex **or** $g\text{-}deposit(\mathcal{CC}_i) \geq k$ **then**
9: mark $\mathcal{CC}_i$ as processed;
10: **for all** $w \in \mathcal{CC}_i$ s.t. $pru(w) =$ false **do**
11: SWEEP$(w, pru, deposit, g\text{-}deposit, \mathcal{CS})$

6. EXPERIMENTS

In this section, we experimentally evaluate the performance of our proposed algorithms.

All algorithms are implemented in C++ using gcc complier at -O3 optimization level. All the experiments are conducted under a Linux operating system running on a machine with an Intel Xeon 3.4GHz CPU, 32GB 1866MHz DDR3-RAM. The time cost of algorithms is measured as the amount of wall-clock time elapsed during program execution.

Datasets. We use 7 publicly available real-world networks to evaluate the algorithms. The network statistics is shown in Table 1.

Stanford is a web graph where vertices represent pages from Stanford University (`stanford.edu`) and edges represent hyperlinks between them. DBLP is a co-authorship network of DBLP. Cnr is a small crawl of the Italian CNR domain. ND is a web graph where vertices represent pages from University of Notre Dame (`nd.edu`) and edges represent hyperlinks between them. Google is a web graph from Google Programming Contest. Youtube is a social network from the video-sharing web site Youtube. Cit is a citation network maintained by National Bureau of Economic Research. All datasets can be downloaded from SNAP[1].

[1] `http://snap.stanford.edu/index.html`

Table 1: NETWORK STATISTICS

Datasets	$\|V\|$	$\|E\|$	Density	Max Degree
Stanford	281,903	2,312,497	8.20	38,625
DBLP	317,080	1,049,866	3.31	343
Cnr	325,557	3,216,152	9.88	18,236
ND	325,729	1,497,134	4.60	10,721
Google	875,713	5,105,039	5.83	6,332
Cit	3,774,768	16,518,948	4.38	793

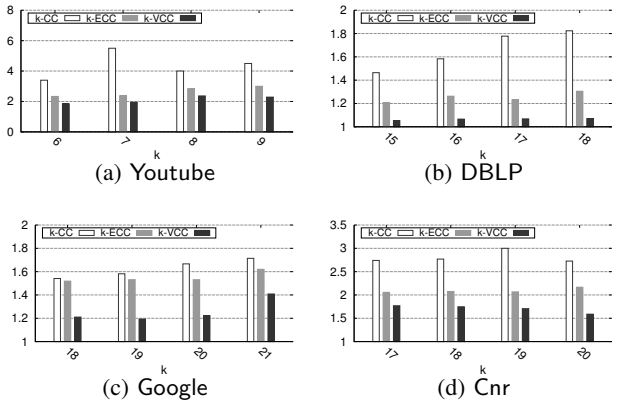

Figure 7: Average Diameter

6.1 Effectiveness Evaluation

We adopt the following three quality measures for effectiveness evaluation:

- *Diameter diam.* The diameter definition is shown in Eq. 1.
- *Edge Density* ρ_e. Edge density is the ratio of the number of edges in a graph to the number of edges in a complete graph with the same set of vertices. Formal equation is given as follows:

$$\rho_e(g) = \frac{2|E(g)|}{|V(g)| \cdot (|V(g)| - 1)} \tag{4}$$

- *Clustering Coefficient* $\mathcal{C}$. The local clustering coefficient $c(u)$ for a vertex u is the ratio of the number of triangles containing u to the number of triples centered at u, which is defined as:

$$c(u) = \frac{|\{(v, w) \in E | v \in N(u), w \in N(u)\}|}{|N(u)| \cdot (|N(u)| - 1)/2} \tag{5}$$

The clustering coefficient of a graph is the average local clustering coefficient of all vertices:

$$\mathcal{C}(G) = \frac{1}{|V|} \sum_{u \in V} c(u) \tag{6}$$

In our effectiveness testings, given a graph G and a parameter k, we calculate the diameter, edge density and clustering coefficient respectively for every k-VCC of G. We show the average value of all k-VCCs for each parameter k. We compute the same statistics for all k-ECCs and k-cores of G as comparisons.

Fig. 7 presents the average diameter of all k-cores, k-ECCs and k-VCCs under the different parameter k in real datasets. Similarly, Fig. 8 and Fig. 9 give the statistics of edge density and clustering coefficient respectively. We choose four datasets Youtube, DBLP, Google, and Cnr as representatives in this experiment. In the experimental results, we can see that for the same parameter k value, k-VCCs have the smallest average diameter, the largest average edge density and the largest clustering coefficient in all three tested metrics. The result shows that our k-VCCs are more cohesive than the k-ECCs and k-cores.

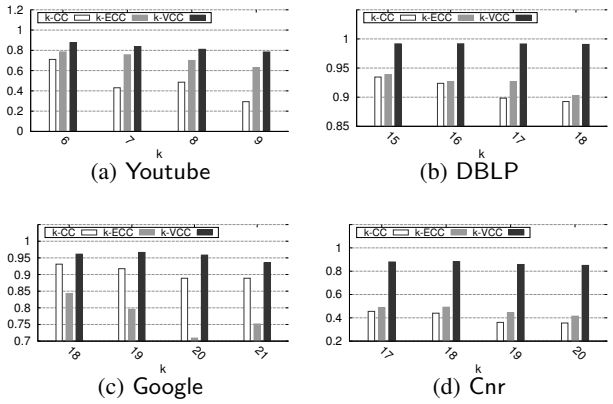

Figure 8: Average Edge Density

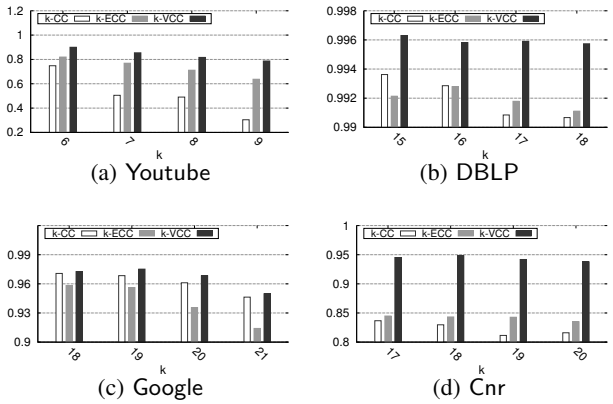

Figure 9: Average Clustering Coefficient

It is worth to mention that in Fig. 7, when k increases, the diameter of the obtained k-cores, k-VCCs, and k-ECCs can either increase or decrease. As an example, the average diameter of k-VCCs decreases slightly while increasing k from 7 to 8 in the Youtube dataset. This is because when k increases, the k-VCCs obtained are more cohesive, thus the diameter for the k-VCCs containing a certain vertex becomes smaller. Such reason also leads to the increase for edge density and clustering coefficient for some k values in these datasets. As another example, the average diameter of k-VCCs increases slightly while increasing k from 20 to 21 in the Google dataset. The reason for this phenomenon is that there exist some small 20-VCCs in which no vertex belongs to any 21-VCC. Here the small k-VCC means there exist small number of vertices inside. These small 20-VCCs have small diameter, which makes the average diameter small. Such reason also leads to the decrease for edge density and clustering coefficient for some k values in these datasets.

A case study is shown in Subsection 6.4 to further demonstrate the effectiveness of k-VCCs.

6.2 Efficiency Evaluation

To test the efficiency of our proposed techniques, we compare the following four algorithms to compute the k-VCCs. For each

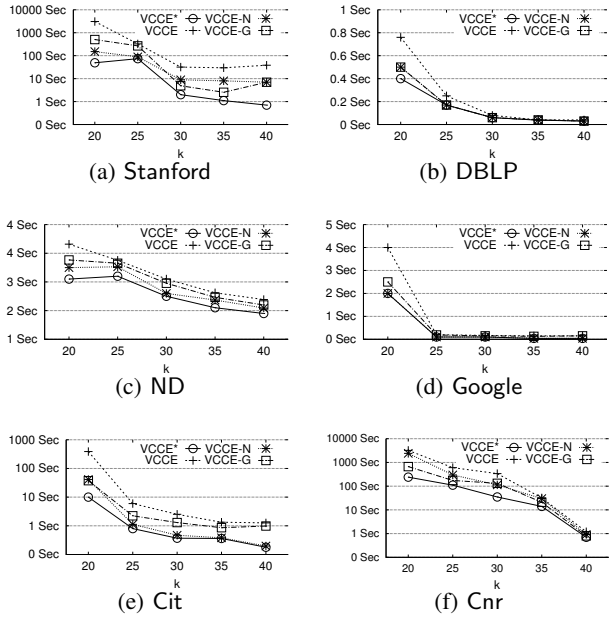

Figure 10: Processing time

Table 2: PROPORTION FOR DIFFERENT RULES

Rules	Stanford	DBLP	ND	Google	Cit	Cnr
NS_1	14%	67%	1%	29%	12%	11%
NS_2	40%	21%	42%	36%	68%	32%
GS	13%	4%	1%	9%	12%	48%
Non-Pru	33%	8%	56%	26%	8%	9%

dataset, we show statistics of algorithms under different parameters k varying from 20 to 40.

- VCCE: Our basic algorithm introduced in Section 4.
- VCCE-N: The basic algorithm with the neighbor sweep strategy introduced in Section 5.1.
- VCCE-G: The basic algorithm with the group sweep strategy introduced in Section 5.2.
- VCCE*: The algorithm with both neighbor sweep and group sweep strategies.

Testing the Time Cost. As we can see from Fig. 10, the time cost of each algorithm generally presents a decreasing trend while parameter k increases. That is because a higher value of parameter k leads to a smaller number of k-VCCs. Intuitively, the algorithm will test less local connectivity during the processing when k increases. A special case here is that algorithm VCCE* spends a little more time under $k = 25$ than under $k = 20$ in the Stanford dataset. This phenomenon happens due to the structure of the Stanford graph in which $k = 25$ leads to more partitions than $k = 20$. We also find that both algorithms VCCE-N and VCCE-G are more efficient than the basic algorithm in all testing cases. Considering the specific structures of different datasets, we find that the group sweep strategy is more effective on graph Cnr, and the neighbor sweep strategy is more effective on other datasets. Our VCCE* algorithm outperforms all other algorithms in all test cases.

Testing the Effectiveness of Sweep Rules. To further investigate the effectiveness of our sweep rules, we also track each processed vertex during the performance of VCCE* and record the number of vertices pruned by each strategy. Specifically, when performing sweep procedure, we separately mark the vertices pruned by *neighbor sweep rule 1* (strong-side vertex), *neighbor sweep rule 2* (neighbor deposit) and group sweep. Here, we divide neighbor sweep into two detailed sub-rules since the both of them perform well and the effectiveness of these two strategies is not very consistent in different datasets. For each vertex v in line 11 of GLOBAL-CUT*, we increase the count for corresponding strategy if v is pruned (line 13). We also record the number of vertices which are non-pruned and really tested (line 12). For each dataset, we record these data under different k from 20 to 40 and obtain the average value. The result is shown in Table 2. *NS_1* and *NS_2* represent *neighbor sweep rule 1* and *neighbor sweep rule 2* respectively. *GS* is group sweep and *Non-Pru* means the proportion of non-pruned vertices. Note that there is a large number of vertices which are pruned in advance by the k-core technique.

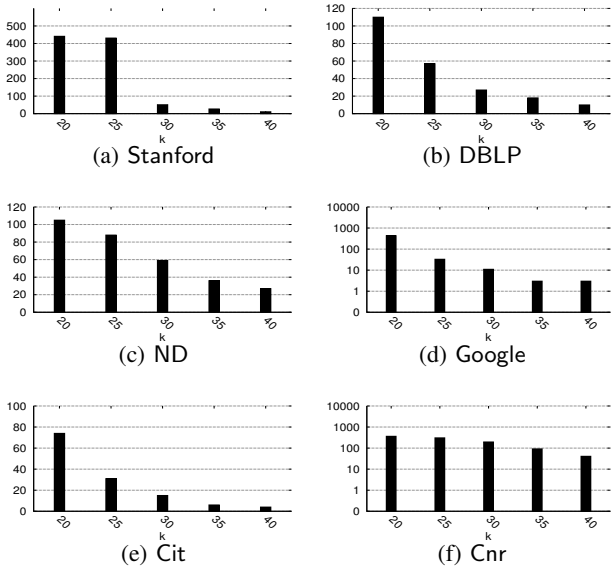

Figure 11: Number of k-VCCs

The result shows our pruning strategies are effective. Over 90% vertices are pruned in DBLP, Cit and Cnr. The proportion of totally pruned vertices is smallest in ND, which is about 45%. Among these pruning strategies, the effectiveness of *neighbor sweep rule 1* and group sweep depends on the specific structure of datasets. *neighbor sweep rule 1* performs much better than group sweep in DBLP. The pruned vertices due to such strategy accounts for 67% of total (including really tested vertices). Group sweep is more effective than *neighbor sweep rule 1* in Cnr. The percentage for group sweep is about 48% while it is only 11% for *neighbor sweep rule 1*. These two strategies are of about the same effectiveness in other datasets. As comparison, the *neighbor sweep rule 2* closely relies on the existing processed vertices. It becomes more and more effective with vertices tested or pruned constantly. Our result shows it is very powerful and stable. The percentage for such strategy reaches to 68% in Cit and is over 20% in all other datasets.

Testing the Number of k-VCCs. The numbers of k-VCCs under different k values for each dataset are given in Fig. 11. The numbers of k-VCCs on all tested datasets have a decreasing trend when varying k from 20 to 40 in Fig. 11. The reason is that when increasing k, some k-VCCs cannot satisfy the requirement and thus will not appear in the result list. The trend of the number of k-VCCs ex-

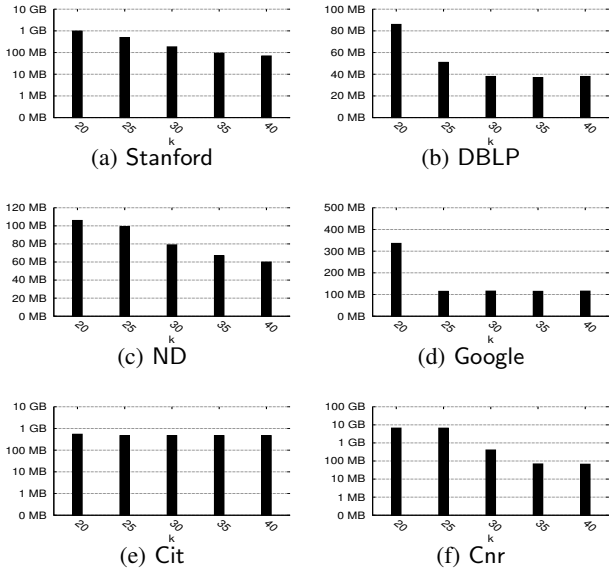

Figure 12: Memory Usage of Algorithm VCCE*

plains why the processing time of our algorithms decreases when k increases in Fig. 10. Note that the number of k-VCCs may vary a lot in different datasets for the same k value. In the same dataset, when k increases, the number of k-VCCs may drop sharply. For example, for when k increases from 20 to 25, the number of k-VCCs in Google decreases by 10 times. The number of k-VCCs depends on the graph structure of each specific graph.

Testing the Memory Usage. Fig. 12 presents the memory usage of algorithm VCCE* on different datasets while verying parameter k. Note that the memory usage of all the four algorithms VCCE, VCCE-N, VCCE-G, and VCCE* are very close since they follow the same framework to cut the graph recursively, and the memory usage mainly depends on the size of graph and the number of partitioned graphs which are the same for all the four algorithms. Therefore, we only show the memory usage of VCCE*.

As we can see from the figure, the memory usage on most of the datasets has a decreasing trend when k increases. The reasons are twofold. First, recall that all vertices with degree less than k are firstly removed for each subgraph during the algorithm. A higher k must lead to more removed vertices, and therefore, makes the graph smaller. Second, when k increases, the number of k-VCCs decreases and the number of partitioned graphs also decreases, which lead to a smaller memory usage. For some cases, the memory usage increases when k increases, this is because when k increases, the sparse certificate of the graph becomes denser, which requires more memory. Generally, the memory usage keeps in a reasonable range in all testing cases.

6.3 Scalability Evaluation

In this section, we test the scalability of our proposed algorithms. We choose two real graph datasets Google and Cit as representatives. For each dataset, we vary the graph size and graph density by randomly sampling vertices and edges respectively from 20% to 100%. When sampling vertices, we get the induced subgraph of the sampled vertices, and when sampling edges, we get the incident vertices of the edges as the vertex set. Here, we only report the processing time. The memory usage is linear to the number of vertices. We compare four algorithms in the experiments. The experimental results are shown in Fig. 13.

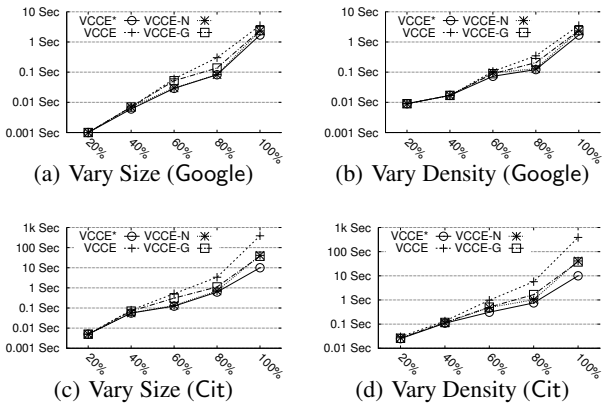

Figure 13: Scalability evaluation

Fig. 13 (a) and (c) report the processing time of our proposed algorithms when varying $|V|$ in Google and Cit respectively. When $|V|$ increases, the processing time for all algorithms increases. VCCE* performs best in all cases and VCCE is the worst one. The curves in Fig. 13 (b) and (d) report the processing time of our algorithms in Google and Cit respectively when varying $|E|$. Similarly, VCCE* is the fastest algorithm in all tested cases. In addition, the gap between VCCE* and VCCE increases when $|E|$ increases. For example, in Cit, the processing time of VCCE* is 20 times faster than that of VCCE when $|E|$ reaches 100%. The result shows that our pruning strategy is effective and our optimized algorithm is more efficient and scalable than the basic algorithm.

6.4 Case Study

In this experiment, we conduct a case study to visually reveal the quality of k-VCCs. We construct a collaboration graph from the DBLP (http://dblp.uni-trier.de/). Each vertex of the graph represents an author and an edge exists between two authors if they have 3 or more common publications. Since the k-VCCs in the original graph are too large to show, we pick up the author 'Jiawei Han' and his neighbors. We use the induced subgraph of these vertices to conduct this case study.

We query all 4-VCCs containing 'Jiawei Han' and the result is shown in Fig. 14 (a). We obtain seven 4-VCCs. Each of them is dense. A vertex is marked black if it appears in more than one 4-VCCs. The result clearly reveals different research groups related to 'Jiawei Han'. Some core authors appear in multiple groups, such as 'Philip S. Yu' and 'Jian Pei'. As a comparison, we get only one 4-ECC, which contains the authors in all 4-VCCs. The result of 4-core is the same as 4-ECC in this experiment. Note that author 'Haixun Wang' appears in 4-ECCs and 4-cores but not in any 4-VCC. That means he has cooperations with some authors in the research groups of 'Jiawei Han', but those authors are from different identified groups and he does not belong to any of these groups.

7. RELATED WORK

Cohesive Subgraph. Efficiently computing cohesive subgraphs, based on a designated metric, has drawn a large number of attentions recently. [5, 7] propose algorithms for maximal clique problem. However, the definition of clique is too strict. For relaxation, some clique-like metrics are proposed. These metrics can be roughly classified into three categories, 1) global cohesiveness, 2) local degree and triangulation, and 3) connectivity cohesiveness.

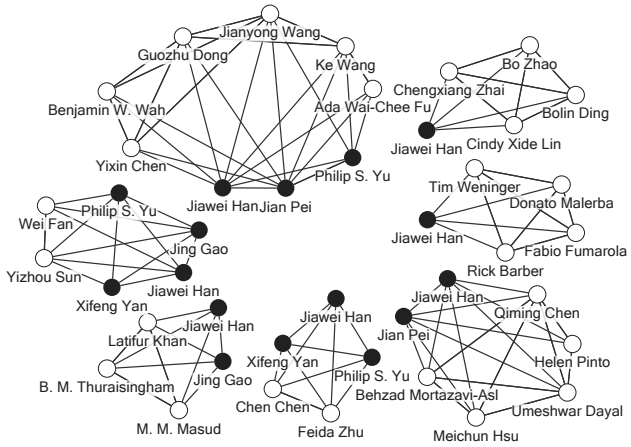

(a) 4-VCCs

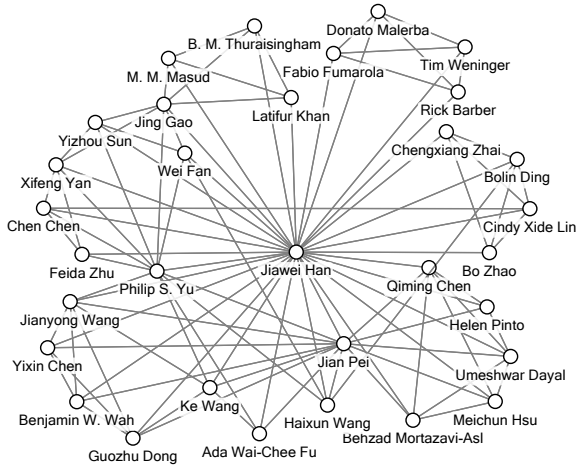

(b) 4-ECCs and 4-core

Figure 14: Case Study on DBLP

1. Global Cohesiveness. [17] defines an s-clique model to relax the clique model by allowing the distance between two vertices to be at most s, i.e., there are at most $s-1$ intermediate vertices in the shortest path. However, it does not require that all intermediate vertices are in the s-clique itself. To handle this problem, [19] proposes an s-club model requiring that all intermediate vertices are in the same s-club. In addition, k-plex allows each vertex in such subgraph can miss at most k neighbors [4, 23]. Quasi-clique is a subgraph with n vertices and at least $\gamma * \binom{n}{2}$ edges [33]. These kinds of metrics globally require the graph to satisfy a designated density or other certain criterions. They do not carefully consider the situation of each vertex and thus cannot effectively reduce the free rider effect [31].

2. Local Degree and Triangulation. k-core is maximal subgraph in which each vertex has a degree at least k [3]. It only requires the minimum number of neighbors for each vertex in the graph to be no smaller than k. Therefore the number of non-neighbors for each vertex can be large. It is difficult to retain the familiarity when the size of a k-core is large. k-truss has also been investigated in [9, 27, 24]. It requires that each edge in a k-truss is contained in at least $k-2$ triangles. k-truss has similar problem as k-core. It is easy to see that two cohesive subgraphs can be simply identified as one k-truss if they share mere one edge. In addition, k-truss is invalid in some popular graphs such as bipartite graphs. This model is also independently defined as k-mutual-friend subgraph and studied in [36]. Based on triangles, *DN*-graph [28] with parameter k is a connected subgraph $G'(V', E')$ satisfying following two conditions: 1) Every connected pair of vertices in G' shares at least λ common neighbors. 2) For any $v \in V \backslash V', \lambda(V' \cup \{v\}) < \lambda$; and for any $v \in V', \lambda(V' \backslash \{v\}) \leq \lambda$. Such metric seems a little strict and generates many redundant results. Also, detecting all *DN*-graphs is NP-Complete. Approximate solutions are given and the time complexity is still high [28].

3. Connectivity Cohesiveness. In this category, most of existing works only consider the edge connectivity of a graph. The edge connectivity of a graph is the minimum number of edges whose removal disconnect the graph. [32] first proposes algorithm to efficiently compute frequent closed k-edge connected subgraphs from a set of data graphs. However, a frequent closed subgraph may not be an induced subgraph. To conquer this problem, [37] gives a cut-based method to compute all k-edge connected components in a graph. To further improve efficiency, [6] proposes a decomposition framework for the same problem and achieves a high speedup in the algorithm.

Vertex Connectivity. [14] proves that the time complexity of computing maximum flow reaches $O(n^{0.5}m)$ in an unweighted directed graph while each vertex inside has either a single edge emanating from it or a single edge entering it. This result is used to test the vertex connectivity of a graph with given k in $O(n^{0.5}m^2)$ time. [13] further reduces the time complexity of such problem to $O(k^3m + knm)$. There are also other solutions for finding the vertex connectivity of a graph [15, 12]. To speed up the computation of vertex connectivity, [8] finds a sparse certificate of k-vertex connectivity which can be obtained by performing scan-first search k times.

8. CONCLUSIONS

Cohesive graph detection is an important graph problem with a wide spectrum of applications. Most of existing models will cause the free rider effect that combines irrelevant subgraphs into one subgraph. In this paper, we first study the problem of detecting all k-vertex connected components in a given graph where the vertex connectivity has been proved as a useful formal definition and measure of cohesion in social groups. This model effectively reduces the free rider effect while retaining many good structural properties such as bounded diameter, high cohesiveness, bounded graph overlapping, and bounded subgraph number. We propose a polynomial time algorithm to enumerate all k-VCCs via a overlapped graph partition framework. We propose several optimization strategies to significantly improve the efficiency of our algorithm. We conduct extensive experiments using seven real datasets to demonstrate the effectiveness and the efficiency of our approach.

9. REFERENCES

[1] J. I. Alvarez-Hamelin, L. Dall'Asta, A. Barrat, and A. Vespignani. k-core decomposition: a tool for the visualization of large scale networks. *CoRR*, abs/cs/0504107, 2005.

[2] G. Bader and C. Hogue. An automated method for finding molecular complexes in large protein interaction networks. *BMC Bioinformatics*, 4(1), 2003.

[3] V. Batagelj and M. Zaversnik. An o(m) algorithm for cores decomposition of networks. *CoRR*, cs.DS/0310049, 2003.

[4] D. Berlowitz, S. Cohen, and B. Kimelfeld. Efficient enumeration of maximal k-plexes. In *Proc. of SIGMOD'15*, 2015.

[5] L. Chang, J. X. Yu, and L. Qin. Fast maximal cliques enumeration in sparse graphs. *Algorithmica*, 66, 2013.

[6] L. Chang, J. X. Yu, L. Qin, X. Lin, C. Liu, and W. Liang. Efficiently computing k-edge connected components via graph decomposition. In *Proc. of SIGMOD'13*, 2013.

[7] J. Cheng, Y. Ke, A. W.-C. Fu, J. X. Yu, and L. Zhu. Finding maximal cliques in massive networks. *ACM Trans. Database Syst.*, 36, 2011.

[8] J. Cheriyan, M. Kao, and R. Thurimella. Scan-first search and sparse certificates: An improved parallel algorithms for k-vertex connectivity. *SIAM J. Comput.*, 22(1):157–174, 1993.

[9] J. Cohen. Trusses: Cohesive subgraphs for social network analysis. *National Security Agency Technical Report*, page 16, 2008.

[10] W. Cui, Y. Xiao, H. Wang, and W. Wang. Local search of communities in large graphs. In *Proc. of SIGMOD'14*, 2014.

[11] J. Edachery, A. Sen, and F. Brandenburg. Graph clustering using distance-k cliques. In *Proc. of GD'12*, 1999.

[12] A. H. Esfahanian and S. Louis Hakimi. On computing the connectivities of graphs and digraphs. *Networks*, 14(2):355–366, 1984.

[13] S. Even. An algorithm for determining whether the connectivity of a graph is at least k. *SIAM J. Comput.*, 4, 1975.

[14] S. Even and R. E. Tarjan. Network flow and testing graph connectivity. *SIAM J. Comput.*, 4, 1975.

[15] Z. Galil. Finding the vertex connectivity of graphs. *SIAM J. Comput.*, 9, 1980.

[16] X. Huang, H. Cheng, L. Qin, W. Tian, and J. X. Yu. Querying k-truss community in large and dynamic graphs. In *Proc. of SIGMOD'14*, 2014.

[17] R. D. Luce. Connectivity and generalized cliques in sociometric group structure. *Psychometrika*, 15, 1950.

[18] K. Menger. Zur allgemeinen kurventheorie. *Fundamenta Mathematicae*, 10(1):96–115, 1927.

[19] R. J. Mokken. Cliques, clubs and clans. *Quality and Quantity*, 13, 1979.

[20] J. Moody and D. R. White. Structural cohesion and embeddedness: A hierarchical conception of social groups. *American Sociological Review*, 68, 2000.

[21] J. W. Moon and L. Moser. On cliques in graphs. *Israel Journal of Mathematics*, 3(1):23–28, 1965.

[22] J. Pattillo, N. Youssef, and S. Butenko. On clique relaxation models in network analysis. *European Journal of Operational Research*, 226, 2013.

[23] S. B. Seidman and B. L. Foster. A graph-theoretic generalization of the clique concept. *Journal of Mathematical Sociology*, 6, 1978.

[24] Y. Shao, L. Chen, and B. Cui. Efficient cohesive subgraphs detection in parallel. In *Proc. of SIGMOD'14*, 2014.

[25] M. Stoer and F. Wagner. A simple min-cut algorithm. *J. ACM*, 44(4), 1997.

[26] A. Verma and S. Butenko. Network clustering via clique relaxations: A community based approach. *Graph Partitioning and Graph Clustering*, 588, 2012.

[27] J. Wang and J. Cheng. Truss decomposition in massive networks. *Proc. VLDB Endow.*, 5, 2012.

[28] N. Wang, J. Zhang, K.-L. Tan, and A. K. H. Tung. On triangulation-based dense neighborhood graph discovery. *Proc. VLDB'10*, 4, 2010.

[29] D. R. White and F. Harary. The cohesiveness of blocks in social networks: Node connectivity and conditional density. *Sociological Methodology*, 31(1):305–359, 2001.

[30] H. Whitney. Congruent graphs and the connectivity of graphs. *American Journal of Mathematics*, 54(1):150–168, 1932.

[31] Y. Wu, R. Jin, J. Li, and X. Zhang. Robust local community detection: On free rider effect and its elimination. *Proc. VLDB Endow.*, 8, 2015.

[32] X. Yan, X. J. Zhou, and J. Han. Mining closed relational graphs with connectivity constraints. In *Proc. of ICDE'05*, 2005.

[33] Z. Zeng, J. Wang, L. Zhou, and G. Karypis. Coherent closed quasi-clique discovery from large dense graph databases. In *Proc. of KDD'06*, 2006.

[34] H. Zhang, H. Zhao, W. Cai, J. Liu, and W. Zhou. Using the k-core decomposition to analyze the static structure of large-scale software systems. *The Journal of Supercomputing*, 53(2), 2010.

[35] Y. Zhang and S. Parthasarathy. Extracting analyzing and visualizing triangle k-core motifs within networks. In *Proc. of ICDE'12*, 2012.

[36] F. Zhao and A. K. Tung. Large scale cohesive subgraphs discovery for social network visual analysis. In *Proc. of the VLDB'12*, volume 6, 2012.

[37] R. Zhou, C. Liu, J. X. Yu, W. Liang, B. Chen, and J. Li. Finding maximal k-edge-connected subgraphs from a large graph. In *Proc. of EDBT'12*, 2012.